"VIRTUAL FIXTURES": PERCEPTUAL OVERLAYS ENHANCE OPERATOR PERFORMANCE IN TELEPRESENCE TASKS

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Louis B. Rosenberg

August 1994

The Virtual Fixtures project was funded by the Air Force Office of Scientific Research (AFOSR) and is therefore entered into the public domain. The research was conducted at Stanford University, Air Force Research Laboratory (Armstrong Lab), and NASA Ames Research Center. The objective of the Virtual Fixtures platform was to research whether it was possible to incorporate interactive virtual objects (called "virtual fixtures") into a user's perception of the real world. This concept later became commonly referred to as augmented reality (AR) or mixed reality (MR). Also see USAF Technical Reports AL/CF-TR-94-0089 and 94-0139.

Virtual Fixtures: Augmented Reality Research (1992)

© Copyright 1994
by
Louis B. Rosenberg

Abstract

Virtual Reality and *Telepresence* systems are human interface platforms which immerse users in environments other than the direct physical surroundings. Operator performance within such *immersive environments* is generally poor compared with the natural surroundings. The typical design strategy for improving the effectiveness of immersive systems has been to increase the *physical accuracy* of the artificial representation. Most telepresence research has focused on increasing the *fidelity* of the link between the operator and the remote site. Most Virtual Reality research has focused increasing the *realism* of display hardware and software. Although enhanced physical accuracy can improve human activity in a simulated reality, such work often leads to ineffecuve or inefficient designs due to misunderstandings about the nature and content of *reality*. The work presented in this dissertation takes a step back, asks basic questions about the nature of reality, and then proceeds with an alternative approach to designing effective immersive systems.

When developing immersive environments, designers often slip into *naive realism* in which they wrongfully equate *reality* with the *physical properties* of their surroundings. This leads to a flawed equivalence between *realistic presentations* and *physically accurate representations*. The design methodology introduced in this dissertation, Design-for-Perception, proposes that reality is a perceptual rather than physical realm. It asks designers to maintain a strict separation in their thinking between the *physical world* and the *perceptual world*. Once designers accept the fact that humans only have access to a highly abstract perceptual representation of their surroundings, they can produce more effective and efficient immersive environments by directly designing for *perceptual reality*.

Application of the Design-for-Perception methodology has lead to an innovative method of enhancing operator performance in telepresence environments. The traditional approach to improving teleoperator performance has been to increase the fidelity (physical accuracy) of the link between operator and remote environment. An alternative solution implied by Design-for-Perception is to enhance operator performance by directly embellishing the user's perceptual reality. *Perceptual-Overlays* are defined as abstract perceptual cues which are introduced into a user's perceptual reality as virtual constructs overlaid on top of the information reflected from the remote site. Because the notion of perceptual overlays is difficult to conceptualize, the *Virtual Fixture* metaphor was formulated to characterize these abstract percepts as concrete entities overlaid on the remote environment. Extensive human testing has shown that Virtual Fixturing can greatly enhance human performance in telepresence tasks. Testing has also shown that telepresence operators can use the virtual fixture metaphor to design their own perceptual overlays and enhance their personal performance in a given task. This serves as a confirmation that perceptual overlays are effective tools for telepresence systems.

Preface

In a nutshell, this thesis is about *immersion*. More specifically, this thesis is about the design and implementation of immersive human interface systems such as virtual reality and telepresence. It asks questions such as what is immersion, how do we achieve it, and how do we use it in an effective and practical application. Although the focus of this thesis is the development of a design methodology for immersive human interfaces, the journey documented in these pages is very similar to that which traveled by students of the visual arts. When learning to draw, a profound revelation often precedes the birth of the artist. That revelation is that although drawing is a seemingly manual skill, learning to draw has little to do with dexterity and a lot to do with perception. Leaning to draw is often described as a process of learning to see. Maurice Grosser writes in *The Painters Eye* "The painter draws with his eyes, not with his hands. Whatever he sees, if he sees clear, he can put down. The putting it down requires perhaps, much care and labor, but no more muscular agility than it takes for him to write his name. Seeing is clearly the important thing." The artist must learn to view a subject, decompose it perceptually, and then reproduce the salient features through some abstract medium. This perceptual reproduction must be such that others can perceive the artist's representation of the original subject.

While this thesis has little to do with learning to draw, it documents a similar journey. This thesis is about learning to become a designer of virtual entities, a process which like drawing requires the student to learn to reproduce a subject through an abstract medium such that others can experience a representation of the original. Like the student learning to draw, I came to a revelation that learning to design virtual environments has less to do with learning to control the hardware and more to do with learning to perceive the world. A design methodology known as *Design for Perception* is proposed in these pages which suggests that virtual percepts be designed though perceptual decomposition rather than physical modeling. Like the artist who reports gaining deeper understanding of visual scenes through drawing, I have gained much insight into physical pheomenon when designing virtual percepts through perceptual decomposition. It is my belief that such levels of understanding could only be gained when a designer investigates physical phenomenon in perceptual rather than mathematical terms.

Chapter One

"The human mind is exquisitely tailored to make sense out of the world. Give it the slightest cue and off it goes, providing explanation, rationalization, understanding"

-Donald Norman
The Design of Everyday Things

---

1. INTRODUCTION

This thesis addresses a wide range of topics including virtual reality, telepresence, human perception, human performance, human-computer interaction, and design. The emerging notion of virtual reality is explored in this work from the perspective of a designer trying to implement this new technology in a useful application. Because there exists no formal methodology for the design and application of virtual constructs, this thesis borrows many ideas from the field of perceptual psychology and attempts to formulate a theoretical foundation for the design of virtual percepts. The fundamental design principle introduced by this thesis, which I refer to as *Design-for-Perception*, is not the only focus of this document, but it permeates all aspects of the design process and is therefor a driving force behind all of the theoretical and experimental work presented here.

At the core of this thesis is the design and implementation of computer generated entities known as *Virtual Fixtures*. This is an application of virtual reality technologies intended to enhance human performance in telepresence tasks by introducing abstract perceptual information into the human-machine interface. Telepresence was chosen as an ideal platform for implementing a concrete application of virtual reality because it shares many of the same hardware requirements for human-computer interaction. Telepresence was also chosen as an ideal platform for the implementation of the principles of Design-for-Perception because *remote presence* involves many of the same perceptual requirements as *virtual presence*. Although the notion of virtual fixturing will be explored in the context of telemanipulation, it is hoped

that this particular implementation will not limit the scope or generality of the results presented here. Many of the same benefits provided by virtual fixtures for telepresence tasks in *remote environments* are likely to hold true for direct manual tasks performed in *natural environments* as well as simulated tasks performed in *virtual environments*.

The remaining pages of this introductory chapter are organized as follows; First, the basic notion virtual reality is presented. Then, basic principles of perceptual psychology are introduced in order to provide a vocabulary and conceptual framework to support the rest of the thesis. Then, the concept of telepresence is presented. And finally, the notion of virtual fixturing is introduced. I apologize for the many levels of introductory text which precede the central notion of virtual fixtures, but a fair amount of background is necessary to properly motivate the fixturing concept.

1.1 From Computer Simulation to Virtual Reality

Emerging technologies for human-computer interfacing are just beginning to provided users with truly *natural* methods of interacting with computational information. Interface technologies such as binocular head-mounted displays, force reflecting masters, tactile stimulators, and binaural audio processors, have evolved to the point where they can replace the standard keyboard-monitor interface as the preferred means of interacting with certain types of information stored within computers. As a result, the traditional approach to computer simulation has suddenly been challenged by a rapidly emerging, poorly defined, and sometimes elusive notion referred to as *virtual reality*. Although still a long way from fruition, the potential of virtual reality has inspired excitement in a broad range of professional communities from science and engineering, to entertainment and education. Virtual reality is often described as the generation of a convincingly realistic, and naturally interactive computer simulation in which the user is given the illusion of being immersed inside of a computer generated environment. The basic premise is that by simulating a physical reality and presenting it to a user through the same sensory channels with which the real physical world is perceived, the resulting *virtual reality* will seem as real to the user as does *physical reality*. If the sensory information is appropriately represented, accurately mapped to the users sensory channels, and correctly modified as the user interacts with the computer model, the user will be given a sense of *presence* within the virtual world in the same way he or she feels present in the real world.

To truly understand the potential of virtual reality, we must look past the recent media hype and ask an obvious, but frequently avoided question; What makes virtual reality different than traditional computer simulation? After all, both generate and present computational models of the real world.

The answer to this question has more to do with the implementation of the simulation than the content. While a traditional computer simulation focuses on generating a realistic *representation* of a physical phenomenon, a virtual reality simulation focuses on generating a realistic *presentation* of a physical phenomenon. Although the difference between representation and presentation may seem subtle, the implications of this shift in perspective can be profound. While users of a traditional simulation are passive observers, the users of a virtual reality simulation are active participants. Within a virtual environment, users can interact with and explore computer simulations by making use of natural sensory and motor skills previously reserved for interaction with the physical world. Given the ability to draw upon their highly evolved human facilities, users of virtual environment systems can obtain an intimate level of insight and understanding of computational models not provided by traditional simulation techniques.

From surgical planning to scientific visualization, many applications of virtual information display have been shown to provide great advantage over traditional interface techniques. A powerful example of the advantages of providing users with an immersive perspective upon computational models is the GROPE system developed by Frederick Brooks [1988] at UNC Chapel Hill. Using a force reflecting master and stereoscopic display, the GROPE systems provides an immersive human interface for chemists to interact with virtual models of molecules. A common task performed by chemists (usually performed with the aid of a traditional computer simulation) is called *molecular docking* in which the chemist must find allowable docking sites between a drug and a protein. This challenging task is a crucial step in the development of an effective drug. The GROPE system replaces traditional simulation techniques by allowing chemists to manipulate virtual molecules by hand actually feel the intermolecular forces when attempting to dock the virtual molecules. As result of this natural and immersive approach to information display, chemists report gaining new insights and understandings of why certain drugs dock well and why others do not. Further more, Brooks reports that the addition of haptic information to the system resulted in a two fold increase in operator performance.

1.2 Design & Virtual Reality, a first pass.

Thus far I have defined virtual reality and addressed the potential benefits of this emerging means of interaction with computational information. As a designer, the next issue to explore is the basic methodology through which one generates virtual constructs and implements these constructs in a useful application. When designing a virtual entity, the ultimate goal is to create a computer simulation which seems so real to the user, the resulting percept could be mistaken for reality. Such a virtual model would be described by a user as looking, sounding, and feeling real. This, of course, does not mean that virtual constructs are limited to simulations of things that actually exist in the real world. It simply means that interaction with a virtual construct, whether it has a real-world analog or not, is such that a user would

accept the simulated entity as being part of the physical world. Thus the ability to produce a suspension of disbelief, at least on a perceptual level, is the basic requirement for the design of an effective virtual construct.

At first perusal, the design problem seems straightforward; simply generate a realistic computer model and present it in such a way that it is consistent with the user's perception of reality. Once such a model has been created and presented, users will be given the feeling of presence within the computer simulation and will thus have access to their natural perceptual and analytical facilities. On the surface, this short description seems to define a simple but effective design methodology for achieving a powerful mode of user interaction with computational information. Unfortunately a designer wanting to pursue the development of virtual constructs will quickly find the above description is far from complete. In fact, the above statements raise deeper questions than they answer, such as; What is *reality*? What constitutes a *realistic* simulation? And finally what is meant by *presence*? As a designer embarking on the challenge of designing and implementing virtual constructs for practical application, it is important to take a brief detour and explore some basic notions of perceptual psychology which make the philosophical issues of reality, realism, and presence more concrete.

1.3 Reality: *Physical or Perceptual?*

What is reality and how does one present information to a user that is so realistic it can be called a virtual reality? That question can only be addressed by answering a further question: What makes a computer generated model seem realistic? The traditional approach to realism in computational simulation has lead many virtual reality researchers to an obvious, but flawed, equivalence between realism and accuracy. In traditional simulations, the goal is to create a mathematical model of a physical phenomenon. Such a mathematical model must be realistic in that it can answer questions and provide quantitative insights into the behaviors of a physical phenomenon. Thus, an accurate representation of the salient physical parameters is essential for computational models used in traditional simulation. When designing a virtual simulation on the other hand, the goal is to present information to a user such that it can be perceived and interacted with as if it were physically real. When generating a model that looks, sounds and feels real to a user, the importance of accurately modeling physical reality is no longer so clear. The change in focus from realistic *representation* to realistic *presentation* requires that designers take a new look at what constitutes a realistic computer model. Nevertheless, designers of virtual entities have a strong tendency to fall into the traditional equivalence between physical accuracy and realism. Interestingly, the reason why so many designers of virtual percepts draw an equivalence between the physical and the realistic is the same reason why they will ultimately find that they do not have to. It stems from a well known phenomenon that perceptual psychologists call *distal attribution*.

Not to single out the designers of virtual percepts as the only people who make the flawed equivalence between physical accuracy and realism, perceptual psychology tells us that *all* people live in a perpetual illusion in which we wrongfully equate our perception of reality with the physical world around us. Referred to by perceptual psychologists as "externalization" or "distal attribution", this mental phenomenon is responsible for our ability to experience daily existence in the physical world despite the fact that we only have access to a highly abstracted perceptual representation of our surroundings. Distal attribution can be defined as follows:

> **Distal Attribution** is the innate human tendency to infer that the mental models of our surroundings, which we construct from abstract perceptual cues, are direct reflections of the physical world rather than the stylized and simplified abstractions of reality which they truly are.

The misconceptions induced by distal attribution are obvious when we consider the fact that we as human beings only transduce a limited band of sensory information with which we construct mental models of the world. Despite our limited bandwidth, we rarely consider our impression of the physical surroundings to be flawed or incomplete. Gestalt psychological theory explains this by suggesting that our mental models of reality are greater than the sum of the perceptual cues from which they are constructed. Thus not only are the mental models we create not a direct reflection of the physical world, they are not even a direct reflection of the abstract perceptual cues which we derive from the physical world. Despite this disparity between the physical world and the mental models we construct, distal attribution is the process which allows us accept our impressions of our surroundings as complete and accurate reflections of reality.

It is within such abstract mental models that we interact with and experience our surroundings. Thus, our internal perceptual model of the world is the only environment which can be truly called *reality*. Despite the fact that we are removed by many levels of perceptual abstraction from the actual physical world, we as human beings experience an undeniable sense of *presence* within our physical surroundings. Loomis [1993] writes "The perceptual world created by our senses and nervous system is so functional a representation of the physical world that most people live their lives without ever suspecting that contact with the physical world is mediate." He goes on to explain that this phenomenon is the result of our natural ability to equate our perceptual reality with physical reality known as distal attribution. Thus, distal attribution is not only responsible for our ability to accept the physical world as reality, it is also responsible for our sense of *presence* within the physical world.

In order to make this discussion or perception, reality, and design more concrete, a number of terms need to be introduced and defined. *Physical reality* refers to the actual world we live in and the physical laws which govern it. This world is not given directly in our experience but is inferred through observation and critical reasoning [Cambell, 1966; Russell, 1948]. *Perceptual reality* , on the other hand,

refers to the mental model of our surroundings which is created based on the information gathered by our senses. This perceptual reality is sometimes referred to by perceptual psychologists as the *phenomenological world.* It is often described as that of which we are perceptually aware, being constructed of our senses and nervous system. It is within this perceptual, or phenomenological, world which the designers of virtual percepts must learn to work. Unfortunately it is a difficult world to find and even harder to hold on to. Loomis [1993] writes that "its high degree of functionality conceals its very nature to the unreflective person." Thus it is important that we as designers maintain a strict separation in our thinking between the phenomenological and physical worlds. [e.g. Brian, 1951; Gogel, 1990; Koffka, 1935; Russell, 1948; Shepard, 1981].

1.4 A Simple Thought Experiment

So slippery is this distinction between physical and perceptual reality, and so powerful is the natural human tendency to make a direct equivalence between the two, I myself often have to stop and re-convince myself that I actually live in a perceptual world rather than a physical world by performing a simple thought experiment. What I generally do is look around my room and ask myself what do I see. No matter where I am, my view can be best described as a world made up of crisp edges. I can trace these edges, follow them with my minds eye. These edges seem as real to me as any other concrete entity in the physical world. But there lies both the problem and the solution. I perceive a physical world made up of edges for the simple reason that my visual system is composed largely of sensors which are designed specifically to respond to edges. Although they seem concrete, edges are completely perceptual entities, being nothing more than a sharp transition between two disparate realms. So low-level is the human tendency to perceive visual reality as a collection edges or transitions, that receptive fields specifically responsive to edges actually exists in the retina, only a few synapses above the raw sensors known as rods and cones [Hartline, 1949]. Long before any information about my surroundings become cognitive, my perceptual reality has been abstracted into a collection of enhanced and exaggerated edges. Of course this process of perceptual abstraction has only just begun with the receptive fields in the retina. Many more levels of abstraction occur as the neurological pathway leaves the retina, make it way to the thalamus, passes through the striate cortex, and finally reaches the cerebral cortex where it becomes the cognitive model we call reality.

Although I will not go into the additional layers of pre-processing which further abstract our visual perception of the physical world before it ever reaches the cognitive centers in the cerebral cortex, I will point out how easy it is to take for granted the many layers of abstraction that actually exist between the physical world and our perception of it. Consider the perception of colors. At an early age we are taught that any color can be composed of a mixture of three primaries. Because of the powerful human

tendency of distal attribution, most people infer that perceptual trichromacy is a property of the physical paints, dies, or lights that are mixed to produce the final color. The fact that any color can be produced by trichromatic combinations seems just as much a property of our physical surroundings as does gravity, inertia, or any other law of nature. Of course there is nothing three dimensional about electromagnetic radiation. The trichromatic perception of color is a purely perceptual rather than physical phenomenon [Helmholtz, 1924]. We as human beings simply have three color receptors with which we decompose the color content of our physical surroundings. Thus any combination of three primaries form a spanning set which can describe the basis of all colors we can perceive. If a person has only two color receptors, we call them color blind because they form a spanning set of their surroundings with only two components and thus they see entire ranges of *different* colors as the same. We use the word *blind* to describe such people because they can not transduce properties of the physical world which are clearly present. Of course egocentric humans like ourselves rarely consider that some other creature having four chromatic receptors would decompose the surroundings into entire ranges of colors which we could not discriminate and would thus consider all unimpaired human beings to be color blind. Although no such creatures are known to exist, this example does point out that the way we perceive our surrounds is based significantly on perceptual effects even though those effects are often incorrectly attributed to physical reality.

A final memorable example of the effect that perceptual abstraction can have upon the interpretation of the physical world can be found by considering the retina of a frog and contrasting it to that of a human. At a similar level of processing where humans have receptive fields sensitive to edges, the frog has receptive fields sensitive to small, dark splotches which move around at a particular frequency [Lettvin, Maturana, McCulloch, & Pitts, 1959]. Sometimes called "bug-detectors", this level of pre-preprocessing allows the frog to abstract physical reality into an environment made up of *bugs* and *not-bugs*. Although such abstraction may seem strange to edge-detecting beings like ourselves, it does in fact produce a very useful and appropriate perceptual reality for such creatures to live in.

1.5 Design-for-Perception

So what do the above examples teach us about the design of virtual percepts? If we were to design a virtual reality for frogs, the requirements for the design of effective virtual constructs would likely be very different than those for humans. Although an obvious conclusion, it does point out that design requirements for virtual percepts are artifacts of perceptual reality, not physical reality (since humans and frogs share the same physical reality). Thus, designers of a virtual environments need not expend the exhaustive modeling effort required to accurately reproduce physical reality. They need only model the highly abstracted perceptual reality and allow the process of distal attribution to foster the equivalence

between the physical with the perceptual. Once this cognitive leap from perceptual reality to physical reality has occurred, the user will accept the computer model as *real* and feel *present* within the simulation. Thus, the basic notion of *Design-for-Perception* simply suggests that perceptual rather than physical parameters be the basis for the design of virtual constructs. Design-for-Perception can be summarized as follows:

> **Design-for-Perception** suggests that perceptual rather than physical paramters be the medium of virtual percept design. More specifically, it states that perceptual reality rather than physical reality is the context in which designs should be born. Furthermore, it suggests that the design of effective virtual constructs only require the representation and presentation of those salient perceptual features which foster the generation of an appropriate perceptual model and allow the process of distal attribution to occur. Distal attribution will instill the desired phyisical realism into the design even if it is not explicitly represented.

Although the above methodology may seem familiar to the visual artist who is trained to abstract the physical world into salient perceptual features through sketching, this notion of *Design-for-Perception* is not nearly as familiar when we explore equally important sensory modalities such as the haptic, auditory, tactile, even vestibular pathways. For example, we as the perceptual designers of virtual constructs must learn not only to *sketch* the visual representation of a physical object, we must learn to sketch the texture, hardness, inertial properties, even thermal properties, such that they will evoke the desired perceptual response in the user. Such an approach to the design of virtual percepts may prove significantly more efficient and effective than the brute force approach of physical modeling. These ideas will be explored in greater detail in Chapter 4.

1.6 Achieving Presence and a Sense of Reality

Having defined the primary design requirement for the generation of effective virtual percepts as the production of an effective perceptual model which fosters the process of distal attribution to equate that model with physical reality, one question still remains unanswered: How does one generate an effective model of perceptual reality such that distal attribution will occur? This is a particularly important question in light of the fact that distal attribution governs our acceptance of a virtual entity as real as well as our illusion of presence within an environment. Although this is an area that is not entirely understood, many ideas have been explored over the last half century. These ideas are best described using some basic vocabulary defined by Von Holst [1954]. He defined "efference" as the commands issued to the muscles of the body. Such commands are very important to our generation of a perceptual reality because they control the relative placement of our sense organs with respect to stimuli. He further defined "afference" as the sensory input which enters our body through our sensory organs. Thus "efference" and "afference" can be

thought of as output and input respectively. White [1970] has proposed that distal attribution occurs in a perceptual environment in which "afference is lawfully related to efference." This can be translated into the insightful statement that the perceptual world seems real when the sensory input remains consistent despite the motions of the body and variations in placement of our sense organs. Loomis [1993] qualified this requirement by stating that not only must there be a systematic covariance between afference and efference for distal attribution to occur, but the observer must be able to identify and understand the relationship in which efference governs changes in afference. By this he means, through training or experience a user must come to understand a predictable relation between motor activity and changes in sensory perception. If such a relation exists but it is too complex for the user to grasp, the effect of distal attribution will not occur. Loomis' qualification upon White's earlier work becomes important when we consider *telepresence systems* in which sensory information from a remote environment is mapped to a user through various levels of distortion.

1.7 Telepresence and the Transparent Interface

The fundamental purpose of a telepresence system is to extend an operator's sensory-motor facilities and problem solving abilities to a remote environment [Pepper, 1984]. Telepresence has been defined by Sheridan [1993] as a human/machine system in which the human operator receives "sufficient information about the teleoperator and the task environment, displayed in a sufficiently natural way, that the operator feels physically present at the remote site." Very similar to virtual reality, in which we strive to achieve the illusion of presence within a computer simulation, telepresence strives to achieve the illusion of presence at a remote location. The end results of both telepresence and virtual reality are essentially the same, a human-computer interface which allows a user to take advantage of natural human abilities when interacting with an environment other than the direct surroundings. Thus, the previous discussion of distal attribution and the generation of presence applies as directly to telepresence as it did to virtual reality. In fact, telepresence is in some ways the ideal platform to explore distal attribution because of the explicit mapping of efferent and afferent channels between environments. The primary difference is that virtual reality requires modeling of the sensory information presented to the user while telepresence is achieved by transducing the actual sensory information from a real remote environment and reflecting that information back to the user. Thus we can look upon telepresence as a prime example of physical modeling in the generation and presentation of a perceptual reality. We shall see that the principles of Design-for-Perception can be applied to supplement and enhance the effectiveness of telepresence systems.

Telepresence is the result of projecting the operator's manipulatory dexterity to a remote environment while reflecting sensory feedback from that environment such that the operator feels present in the remote site [Akin, Minsky, Thiel, and Kurtzman, 1983]. Such systems are usually achieved through

the use of a vast array of robotic equipment and human interface tools. The hardware generally includes: a highly sensored remote manipulator capable of tracking the operators motions and reflecting sensory information back the operator space, a communication link between the remote manipulator and the operator space, and a user interface capable of naturally presenting the information gathered in the remote site to the operator as well as relaying the operators kinesthetic activities to the remote manipulator. Describing such a configuration in terms of the vocabulary introduced by Von Holst [1954], a telepresence system is one in which a user's *efferent* motor commands are mapped to a remote manipulator while the *afferent* sensory input is reflected from that remote manipulator. For a user to achieve a sense of presence in the remote site, a predictable and understandable covariance between afference and efference must exist [White, 1970; Loomis, 1993]. For example, when the operator moves his or her arm to grab an object in the remote environment, the remote manipulator must follow such kinesthetic activities and reflect the appropriate haptic cues when the object is encountered. If such mappings are incorrect, incomplete, or too complex to grasp, the illusion of presence in the remote environment will be greatly reduced.

Despite the fact that telepresence systems have direct access to the physical phenomenon in the remote environment being represented, limitations in human interface hardware and communication links often result in an incomplete, and often spatially and temporally distorted representations of perceptual reality. As a result, the illusion of remote presence is poorly achieved if achieved at all. Since the goal of achieving presence is to enhance performance in remote environments by allowing operators to make use of their natural physical and mental facilities at the remote site, degradation in presence directly results in degradation in performance. Thus a major goal of most research into telepresence is to find means of enhancing operator performance within remote locations.

To enhance operator performance and understanding within remote environments, most research and development of telepresence systems has been directed towards increasing the *transparency* of the link between operator and environment. Much of this work has focused upon improving the fidelity and presentation of the reflected sensory information. Other work has focused upon making the *projection* of the operator's kinematics and dynamics more naturally mapped into the remote environment [Burdea, 1991; Siva, 1990; Julian and Anderson, 1988; Jacobsen et al., 1991]. The culmination of such research efforts could be a transparent link between human and machine; a user interface through which information is passed so naturally between operator and environment that the user achieves a complete sense of presence within the remote site.

Although unattainable by current technology, the linking of an operator to a remote environment through a perfectly transparent interface should yield a human-machine system that allows an operator to perform just *as well* in the remote environment and in a real environment. Although the achievement of

such an interface would be an important milestone in human-machine systems, the motivation for the work described in this thesis is very different. Rather than looking for methods of enhancing the transparency of the interface, this thesis is actually focused upon corrupting the transparency of the interface by introducing abstract perceptual information between human and machine. By applying the principle of Design-for-Perception, and recognizing that human beings function in a perceptual rather than physical representation of the world, this thesis attempts to enhance operator performance in telepresence systems by embellishing perceptual reality rather than enhancing the representation of physical reality. It is believed that abstract percepts overlaid on top of the reflection of the workspace can be implemented as perceptual tools that enhance an operator's performance and understanding. It is not suggested that such an approach is a replacement for a high fidelity link between operator and environment, but rather a supplement. The culmination of work along these lines could be a telepresence system which does more than simply *project* an operator's abilities into a workspace, but actually *enhances and expands* the operator's abilities within that workspace beyond natural capabilities.

1.8 The Concept of Perceptual Overlays

I define a *perceptual overlay* as abstract sensory information overlaid on top of the reflected sensory feedback from a virtual, remote, or real environment. Since perceptual overlays are laid on top of environments in which an operator feels present, I will refer to the virtual, remote, or real environment on which perceptual overlays are laid as the *ambient reality*. For the particular applications addressed in this thesis, the ambient reality will be the reflected perception of a remote telepresence worksite in which a task is to be performed. A perceptual overlay may be composed of sensory information from a wide variety of modalities including visual, auditory, haptic, and tactile sensations. I do not want to impose any restrictions upon the nature or content of a perceptual overlay, other than to require that the information presented in a perceptual overlay is such that a natural and predictable relation exists between an operator's neuromotor activities (efference) and the subsequent changes in the sensations included in the perceptual overlay (afference). This single restriction upon the nature of perceptual overlays is sufficient to encourage distal attribution to occur [White, 1954; Loomis, 1993]. As a result of distal attribution, the operator will accept the overlaid percepts as real and tangible parts of the ambient reality. In other words, the operator will achieve a sense of presence with respect to the overlaid perceptual information. Thus the virtual percepts will be accepted as genuine features or properties of the environment in which the operator is working. In describing effective modes of feedback from user interfaces, Donald Norman writes "feedback must provide information that matches the users intentions and must be in a form that is easy to understand" [Norman, 1988]. In light of this requirement, perceptual overlays are an ideal form of feedback, providing information about a user's activities which are in an inherently expected and

understandable form. Since perceptual overlays are in fact virtual constructs, and not truly a part of the distal environment, perceptual overlays can be seen as embellishments of the user's perceptual reality. Thus the construction of perceptual overlays lends itself to the implementation of the *Design-for-Perception* methodology because perceptual overlays are entities which only exist in an operator's perceptual reality, having no physical analog to be attributed to.

The particular application of perceptual overlays addressed in this thesis is a use in telepresence systems to enhance operator performance in manual tasks executed within remote environments. This work reveals that by overlaying particular configurations of perceptual information on top of the reflected sensory feedback from a telepresence worksite, operator performance can be greatly enhanced in a variety of manual tasks. To convey the implications of perceptual overlays on human performance in manual tasks, it is convenient to provide the following real world analogy. Consider this: when asked to draw a straight line in the real world, human performance can be greatly enhanced by using a simple tool such as a ruler. The use of a ruler reduces the amount of mental processing required to perform the task, speeds an operator's line drawing ability, and most of all allows an operator to draw a significantly better line than if no ruler had been used at all. Without a ruler, line drawing is a manual task that requires constant visual supervision and hand/eye coordination. With a ruler, line drawing is not only faster and straighter, but the dependence on visual feedback is reduced, freeing up that modality for other uses. What is more, a ruler is often used as a barrier to protect against dangerous or destructive failures, protecting the work-piece from the slip of a pencil or knife. Such guidance and protection allows the operator to ease mental criteria for task success and failure, reducing the level of concentration devoted to the task. Although a simple tool by any standard, a common ruler is clearly a powerful performance aid in manual line drawing tasks.

Although the use of a ruler to assist in straight-line drawing is an effective means of enhancing human performance, can such a process be generalized beyond line drawing? Ruler-use can be seen as a process of *overlaying* abstract perceptual information on top of a workspace. Thus, a ruler can be thought of as a particular perceptual overlay designed to enhance line drawing performance. In the case of a ruler, the overlaid sensory information represents a rigid surface that is perceived haptically and visually by the user. By overlaying this generic piece of perceptual information on top of a workspace, the user has reduced the mental and physical demands of the straight line drawing task and performance is greatly enhanced. If a simple ruler-like perceptual overlay can so greatly enhance the performance of real world manipulatory tasks like straight line drawing, it seems that computer generated perceptual overlays could be developed within virtual environments to enhance the performance of telemanipulation tasks within remote worksites. Just as a ruler can be overlaid on top of a real workspace, such virtual perceptual overlays could be overlaid on top of the sensory feedback from a remote workspace.

1.9 The Virtual Fixture Metaphor

Because the notion of abstract perceptual information being overlaid on top of an ambient reality is as difficult to conceptualize as it is to talk about, I will introduce the *virtual fixture metaphor* as a means of describing perceptual overlays as concrete physical structures. The purpose of the virtual fixture metaphor is to facilitate the understanding of, and interaction with perceptual overlays. The fixture metaphor provides users with a familiar *conceptual model* with which to internalize the notion of perceptual overlays. In psychological terms, the virtual fixture metaphor is intended to facilitate the acceptance of perceptual overlays as part of an ambient reality by providing a context of distal physical entities for the perceptual cues to be attributed to. Without this metaphor, the process of distal attribution would still occur, but it would be unclear to the user which distal stimulus was responsible for the perceived virtual percepts. With this metaphor, the ambiguity is resolved by allowing all users to attribute perceptual overlays to concrete distal stimuli known as virtual fixtures. Regardless of the intent of the virtual fixture metaphor, I would like to stress that the metaphor should not be taken so literally as to limit the scope or application of the perceptual overlay concept. *Virtual fixtures* are hereby defined as follows:

> **Virtual Fixtures** are abstract percepts overlaid on top of the reflected sensory feedback from a remote environment such that a natural and predictable relation exists between an operator's kinesthetic activities (efference) and the subsequent changes in the sensations presented (afference).

Thus virtual fixtures are computer generated percepts overlaid on top of the sensory feedback from a remote telepresence worksite. The operator of a telepresence system should simultaneously feel *present* with respect to both the *real* remote environment and the *virtual* overlays. Although overlaid on top of the user's perception of the remote environment, virtual fixtures are completely independent of all information from the remote site (because they are embellishments of the operator's perceptual reality and are not related to the physical reality of the remote environment). In other words, the perceptual information presented to a user as a result of interaction with a virtual fixture only depends on the computational model which defines the fixture and the kinesthetic activities of the user. Thus virtual fixtures for telepresence systems are independent of all feedback from the remote environment, making them immune from the communication delays and bandwidth limitations which distort and disrupt the perceptual information reflected from the remote site. Thus in many ways, virtual fixtures are actually the highest fidelity representations within the telepresence operator's perceptual reality.

Like the ruler guiding the pencil, virtual fixtures overlaid on top of a remote workspace can reduce mental processing required to perform the task, reduce the work load of certain sensory modalities, and most of all allow precision and performance to exceed natural human abilities. Although virtual fixtures could be functionally equivalent to fixtures in the real world, there are many advantages inherent to virtual fixtures because they are computer simulations rather than real physical hardware. When overlaid on a workspace, the fixtures only interact with the user and not with the workspace itself. Thus fixtures can occupy the same physical space as objects in the remote environment. This means that the workspace geometry imposes no constraints upon the placement or configuration of virtual fixtures. What is more, such fixturing has no mass, has no physical or mechanical constraints, requires no machining time or maintenance, can be easily prototyped and modified, and can essentially be transported to remote locations using nothing more than standard communication links.

If we explore the concept of virtual fixtures using the simple ruler example as the starting point, the first elements to consider might be rigid planar surfaces. Such fixtures would be composed of haptic sensations generated by reflecting simulated forces to the operator through a force-reflecting master. As the operator interacts with the modeled surfaces, the reaction forces would be computed and reflected appropriately. Of course, such fixtures are by no means limited to rigid surfaces. Abstracting the fixturing concept, we might consider modeling compliant surfaces, damped surfaces, even viscous or coulomb frictional contacts. In fact, a virtual simulation environment offers such freedom that fixtures could even be developed as attractive or repulsive fields. Although fixtures composed of haptic sensations offer endless possibilities, the fixturing concept is not limited to that modality. Abstract fixtures could be composed of visual, auditory, even tactile sensations used alone or in cross-modal combinations. For example, if a haptic fixture composed of rigid planar surfaces or attractive force fields was developed to aid a particular task, an audio, visual, or vibratory signal could be mapped to various locations along the fixture to enhance interaction. Such additional modalities could be used to indicate deviations from a trajectory, proximity to a danger zone, even provide feedback of velocities or accelerations.

Exploring the fixturing concept further, we could imagine virtual fixtures imbued with particular visual qualities to enhance interaction with environments. For example, virtual fixtures composed of haptic surfaces could be modeled with optical properties to suit the task at hand. The fixture might be invisible to the user if the operator gains no benefit from visual cues, it could be made to look like a solid object if rich visual cues are useful for the task, it could even be made to look like a transparent glassy solid if visual cues are important but the user wants to avoid obscuring the workspace. Fixtures might even be designed as visual filters to block particular distractions, enhance contrast, provide depth cues, even magnify a part of the workspace. One can even imagine the benefit of a fixture composed of compliant surfaces which changes color or brightness with compression. Previous work with targeting cues

[T. Books, et al., 1991], predictive displays [Bejczy et al., 1990; Oxenberg et al., 1988; Russell et al., 1986; Sheridan, 1986; Sheridan and Ferrell, 1981], and perspective overlays [Kim et al., 1987; Stark, et al. 1989] has demonstrated that overlaid visual cues can enhance performance and understanding within a teleoperation environment.

If the description of virtual fixtures thus far seems too abstract, a simple example may drive the concept home. Imagine a situation where a teleoperating surgeon performs a delicate procedure on a patient. Although such uses of telerobotics are still in the research phase, it is an application that demands a high degree of human performance within a remote workspace. Now imagine that a virtual fixture is being used by the doctor to enhance his abilities in this procedure. The fixture might appear to the doctor like a flat plane of glass with a grooved guide for the scalpel. The glass-like virtual fixture might actually pass directly through a patient's body, preventing the scalpel from penetrating below a particular depth but not obscuring vision of the tissue below. By sliding the scalpel along the edge of a groove in the fixture, the surgeon could make a perfect incision. The slightest deviation from the target trajectory might be reported by an audio or tactile signal. The power and flexibility of such a system would be unmatched by actual physical tooling. Besides the fact that such a fixture in the real world could not pass directly through a patient's body, it could not be put in place at the touch of a button, removed at the touch of another button, or easily altered as conditions change. What is more, virtual fixturing does not have to be fabricated, sterilized, transported, or maintained.

Although the above example of virtual fixture usage is convincing, it often gives readers the misconception that a virtual fixture needs to be carefully designed, highly precise, and accurately calibrated to a remote environment in order to be an effective tool. Although precise virtual fixtures do have potential applications, the virtual fixture concept is in no way limited by the requirement for precise models, accurate hardware, and careful calibration. To dispel this misconception, I will propose an additional example in which a loosely defined virtual fixture could be used effectively to enhance operator performance. Imagine a telepresence system intended for use in the clean-up a hazardous environment such as laboratory contaminated with radioactive material. Such clean-up applications are generally cited as prime applications for telepresence systems. Now imagine that an operator is controlling a remote manipulator arm to pick up the radioactive waste scattered throughout the worksite and deposits the waste in a special receptacle. Finally, imagine that there is an obstacle in the workspace such as a water-main or natural gas line which needs to be avoided by the remote manipulator throughout the clean-up task. One can easily envision a virtual fixture which is simply a repulsive field which provides a haptic perceptual cue to the operator generated as resisting force proportional to proximity to the obstacle. Or maybe a haptic fixture is simply a flat rigid surface placed a few feet in front of the obstacle such that it would be encountered before the manipulator came within striking distance of the obstacle. Although such haptic

perceptual overlays need not be particularly precise, they could enhance operator performance in a number of ways. First, by providing haptic cues which indicate proximity to the obstacle, the operator's visual modality could be freed up for other uses such as scanning the worksite for more targets to pick up. Second, by providing resistance or even a physical barrier against encountering an obstacle, such fixtures could reduce the mental processing required by the task, reduce the criteria for successful task completion, and enhance operator confidence during task execution.

1.10 Perceptual Workstation Environment

Although the notion of virtual fixturing does seem promising, if the development of effective fixtures requires complex computation or intimate prior knowledge of the remote workspace, robot autonomy might be a preferred solution to the task at hand. If, on the other hand, effective fixtures could be rapidly developed by the human operator of a telepresence system through the use of an interactive fixture design interface, virtual fixtures could be used in unstructured or changing environments unsuitable for autonomous systems. What is more, if we consider that the basic goal of telepresence systems is to extend an operator's natural manual and cognitive abilities to a remote environment, it seems that providing the ability to interactively design and modify assistive virtual fixtures is essential to preserving the flexibility of the human in the loop. Thus, it is proposed that a workstation type computer environment be developed to allow operators of telepresence systems to design and implement assistive fixtures when confronted with a unknown task in an unstructured environment. Such a workstation would provide users with a *toolbox* of basic fixture elements from which effective virtual fixtures could be quickly composed to tackle a given task. With such a workstation in mind, the studies described in this thesis investigate the effectiveness of virtual fixtures composed of very simple haptic and auditory percepts to evaluate if effective fixtures can be constructed simple building blocks. Although a workstation environment could implement more sophisticated percepts including a more diverse array of sensory modalities, it was thought that if simple combinations of forces and sounds could be made into effective fixtures, the potential of virtual fixtures would be adequately displayed and the viability of an interactive perceptual workstation would be supported.

Of course there are a few final issues which must be addressed in order for the notion of an interactive perceptual workstation for virtual fixture design to be adequately supported: namely, can users actually design and implement their own virtual fixtures for a given task. If so, can they learn to design such fixtures without extensive training and guidance? More importantly, can they design fixtures that are effective in significantly enhancing manual performance in a given telemanipulation task? If such a phenomenon can be shown, the true potential for an interactive fixture design workstation for telepresence

systems would be revealed. Thus, it is proposed that a limited prototype of such a perceptual workstation be developed and tested on human subjects as the final confirmation of the viability and potential of this novel application of virtual reality.

1.11 A Field Guide to the rest of this Document

The remaining chapters of this thesis will be organized as follows: Each chapter will sequentially present a deeper exploration into the notion of *Design-for-Perception* through empirical testing of the *virtual fixture* concept. Each of the chapters 2 through 7 are self contained units, each presenting a complete and independent description of a series of human experiments. Thus each chapter of this thesis can be taken out of context and read for its own empirical and theoretical value. Nevertheless, I suggest the reader progress through each of these experiments in the order presented in this document because they are organized in such a way that each subsequent set of experiments attempts to address some of the major questions raised by previous experiments. Although many levels of experimentation are documented in these pages, the reader should not take this to mean that this thesis is intended to be a comprehensive exploration of all facets of Design-for-Perception and its implications on virtual fixturing. Although this document does presents detailed empirical results, this work only reveals the potential of Design-for-Perception and virtual fixturing and actually raises many more questions than it answers. The last chapter of this document attempts to put this work in the context of a larger exploration into the ideas introduced by this thesis and provides suggestions for future study. The following section is a listing of all chapters in this document with a brief description of the work presented within.

1.12 Brief Chapter Overview

Chapter 1 Introduction:

-This chapter introduces Design-for-Perception, Perceptual Overlays, and Virtual Fixtures.

Chapter 2 Preliminary Investigation of Perceptual Overlays:

-This chapter provides the first empirical results demonstrating that virtual fixtures can enhance operator performance. Fixtures are not computer generated in this proof of concept study.

Chapter 3 Virtual Fixtures and Time-Delay Telepresence

-This chapter extends the testing of virtual fixtures to the case when time delays exist between the operator and the remote environment. Fixtures are not computer generated in this preliminary study.

Chapter 4 Design-for-Perception and the Rigid Wall Percept.

-This chapter takes a detour from the testing of virtual fixtures and explores the more general concept of the design of virtual percepts. This chapter addresses issues related to the design of computer generated sensations so that future studies can implement computer generated fixtures. The notion of *perceptual decomposition* is introduced and tested empirically.

Chapter 5 Computer Generated Haptic Fixtures in 3D

-This chapter is the empirical testing of computer generated haptic perceptual overlays. This chapter also extends the testing of virtual fixtures to three dimensional space.

Chapter 6 Binaural Perceptual Overlays: Audio Fixtures in 3D Space

-This chapter empirically tests the use of 3D audio fixtures composed of abstract binaural localization cues.

Chapter 7 The Perceptual Workstation

-This chapter brings the notion of virtual fixtures out of the realm of controlled psychophysical testing and into the realm of practical application. A prototype *perceptual workstation* was developed which allows subjects to perform a complex, simulated telemanipulation task and implement virtual fixtures. This is the first set of empirical tests in which subjects are required to *design their own* virtual fixtures for a given task using a *tool box* of perceptual elements.

Chapter 8 Where we are and where we are going

-This chapter sums up the results of the experimental work presented in this thesis and makes recommendations for future studies.

1.13 Glossary of Terms:

Afference: The sensory input which enters our body through our sensory organs

Ambient Reality: The direct perceptual representation of reality which incident upon a person. This could be the direct physical surroundings, a remote telepresence worksite, or a virtual environment.

Design-for-Perception: A design methodology for the construction of virtual percepts which suggests that effective and efficient virtual constructs can be developed through directly modeling the perceptual content of an object rather than physically modeling the object itself. The key concept is that the design of effective virtual constructs only requires the representation and presentation of those salient perceptual features which foster the generation of an appropriate perceptual model and allow the process of distal attribution to occur.

Distal Attribution: A psychological phenomenon which causes humans to naturally make a flawed equivalence between the perception of our environment and the physical world around us. It is this process which allows us to accept the highly abstracted and stylized perceptual representation of our surroundings as a direct reflection of physical reality.

Distal Stimulus: A distant source of sensory information.

Efference: The commands issued to the muscles of our body

Exoskeleton: A human interface which is worn as an external structure which tracks human limb position and often reflects forces back to human limbs.

Fitts' Law: A standard human performance measurement battery in which operator performance in manual tasks is quantified in terms of information processing capacity. The basic premise is that the maximum speed at which an operator can perform a given task is bounded by the neuromuscular systems ability to control movements. Typical Fitts' Law tests involve simple point-tapping or peg-insertion paradigms.

Force Reflection: The act of presenting computer generated force information to a human limb.

Haptic: Sensory modality related to the unified and coordinated activity of tactile and kinesthetic senses. As described by V. B. Brooks, the haptic sense is equivalent to "active touch" and is involved in hand based explorations of physical objects. In this thesis, as well as in many papers in the field, haptic is used as a general term to refer to human force perception.

Haptic Display: Although the term haptic refers to both force and tactile perception, the term *haptic display* generally refers to just the display of force information and is a distinct from tactile display. Thus a *haptic display device* is a piece of force reflecting hardware such as an actuated joystick or exoskeleton.

Kinesthesia: A sensory modality related to the sense of limb position or limb motion in space. It is perceptual system which provides humans with an innate sense of their body position.

Percept: While *perception* is a process, a *percept* is the result of that process. A percept is a mental model of some aspect of the outside world which is created based upon the information gathered by our senses.

Perceptual Decomposition: A perceptual analysis done to reveal the salient perceptual features which compose a given percept. The ultimate goal of such a decomposition is to reveal what features should be included in an effective and efficient virtual representation of the given percept.

Perceptual Hypothesis: When provided with a proximal stimulus, the act of perception is often described as an act of guessing what distal stimulus must have caused it. This act of inference in the process of perception is called the Perceptual Hypothesis.

Perceptual Overlay: Abstract sensory information which is overlaid on top of the reflected sensory feedback from a virtual, remote, or real environment. In the context of this thesis, perceptual overlays are assumed to be computer generated constructs.

Perceptual Reality: The mental model of our surroundings which is created based on the information gathered through our senses.

Physical Reality: The actual world we live in and the physical laws which govern it.

Presence: A state of accepting the ambient sensory information as the physical surroundings. For virtual presence and telepresence this involves a suspension-of-disbelief because the operator **knows** that the ambient reality is not truly the direct physical surroundings.

Proprioception: Is the sensory modality related to a person's innate sense of limb dynamics.

Proximal Stimulus: The sensory information falling upon your receptors.

Remote Environment: A physical site whose properties are sensed and feed to a user in a location other than the sensed site.

Remote Manipulator: A robot placed in a remote environment whose actions are under the control of an operator in another location.

Telepresence: A human interface system between an operator and remote environment in which information is passed so naturally that the human operator actually feeis present in the remote location. Sensory feedback is reflected from the remote environment to the human operator while at the same time, the operator's motions are mapped into the remote environment though robotic devices at that location.

Telemanipulation: Activities of a robot in a remote site under the control of a human operator in a distant location.

Virtual Environment: A simulated space in which is presented to a user through a Virtual Reality interface such that the user feels immersed or *present* within the simulation.

Virtual Fixtures: Computer generated perceptual information overlaid on top of the reflected sensory feedback from a remote environment such that a natural and predictable relation exists between an operator's kinesthetic activities and the subsequent changes in the sensations presented. The intent of Virtual Fixtures is to enhance operator performance and understanding in remote and virtual environments.

Virtual Reality: A computer generated environment presented to a user such that the user feels immersed or *present* within the simulation.

Chapter Two

"The properties of the human operator is particularly important in situations in which the human rather than the manufactured elements provide the essential limitations of the system's performance."

-Jackson Beatty
Handbook of Perception and Human Performance

---

2. Preliminary Investigation of Perceptual Overlays

This chapter explores the viability of the design for perception methodology by experimentally testing the notion of *perceptual overlays* in an actual telemanipulation task. This preliminary set of tests was intended to reveal whether or not embellishments of a user's perceptual reality, in the form of virtual fixtures, can be effective in enhancing manual performance in a controlled telepresence paradigm. Virtual fixtures were implemented in these tests as combinations of two dimensional haptic and auditory percepts overlaid on top of the reflection of a remote telepresence environment. Although these studies were limited to haptic and auditory modalities and were restricted to two dimensions, the results confirmed that perceptual overlays in the form of virtual fixtures can enhance operator performance in telepresence tasks. Just as tools and fixtures in the real world can enhance human performance by guiding manual operations, providing localizing references, and reducing the mental processing required to perform a task, virtual fixtures, overlaid on a remote workspace were found to provide similar benefits.

Eight human subjects participated in this round of testing. They performed telepresence peg-insertion tasks with and without the aid of virtual fixturing. A Fitts' law paradigm was used to quantify operator performance for a variety of virtual fixture configurations. It was found that virtual fixtures composed of haptic and auditory perceptual overlays could increase operator performance up to 70% over the baseline no-fixture trials. The following pages describe the Fitt's Law telemanipulation task as well as the experimental hardware and testing procedure which allowed the implementation of perceptual overlays in

an actual telepresence task. Finally, this chapter presents the results of subject testing and reviews the conclusions which such testing revealed.

2.1 Human Performance Measurements

To quantify teleoperator performance in a remote manual task, a Fitts' Law paradigm was chosen because of its general acceptance as a robust measure of human performance [Fitts, 1954]]. Although extensive use of Fitts' Law has been documented in human performance literature, little work has been done to extend the paradigm to a telepresence environment. McGovern [1975] used a Fitts' Law task to demonstrate the merit of a closed loop master-slave system as compared to an open loop system. Hill [1979] used a Fitts' task to demonstrate performance differences as a function of force feedback from a manipulative system. Pepper [1988] was the first to use the Fitts' task in a true telepresence scenario. While previous work had employed Fitts' Law with the teleoperator in direct view of the workspace, this work used Fitts' Law to compare a variety of viewing conditions which included remote visual links between operator and workspace. Cannon [1992] derived Fitts' law non-empirically using automatic control theory to predict human-machine performance of a teleoperator and confirmed the applicability of Fitts' Law to such systems. All of these previous studies have shown that a Fitts' Law paradigm is appropriate for analysis of perceptual-motor performance within teleoperated and true telepresence systems.

Fitts [1954] established a means of quantifying human performance in terms of information processing capacity of the neuromotor system. He developed a relationship between the speed and the accuracy of human motor performance and demonstrated that the speed of a task requiring a particular accuracy is bounded by the capacity of the neuromuscular system to control movements. Fitts argued that if manual control was limited by the information processing rate of the peripheral and central nervous system, movement times would be limited by the information processing demands of the task. This concept is apparent if we think of a task requiring great accuracy such as threading a needle. Why do we perform a threading task with infuriating sluggishness? Such a task is limited by human information processing capacity (i.e., how fast we can perceive the environment, actuate our limbs, and adjust for error). A task such as threading a needle requires many fine adjustments and is thus limited by how quickly perceptions and adjustments can be performed (i.e., limited by human bandwidth).

To measure human information processing capacity as an indication of performance within a telepresence environment, subjects were tested on a peg insertion task. Like threading a needle, the difficulty of peg insertion is a function of tolerance, the difference between the peg diameter and the hole diameter. The tighter the fit of the peg in the hole, the more accuracy required of the operator, and thus the

slower the maximum performance speed. By requiring subjects to perform a standardized peg insertion task as fast as they can, completion times could be measured and compared to task difficulty to yield information processing capacity for that task. By doing such an analysis upon subjects with and without the implementation of various virtual fixtures, variations in task completion times indicate performance changes resulting from fixture use. Thus by using Fitts' Law analysis of a standardized peg insertion task, the effectiveness of various virtual fixtures as perceptual aids could be quantified.

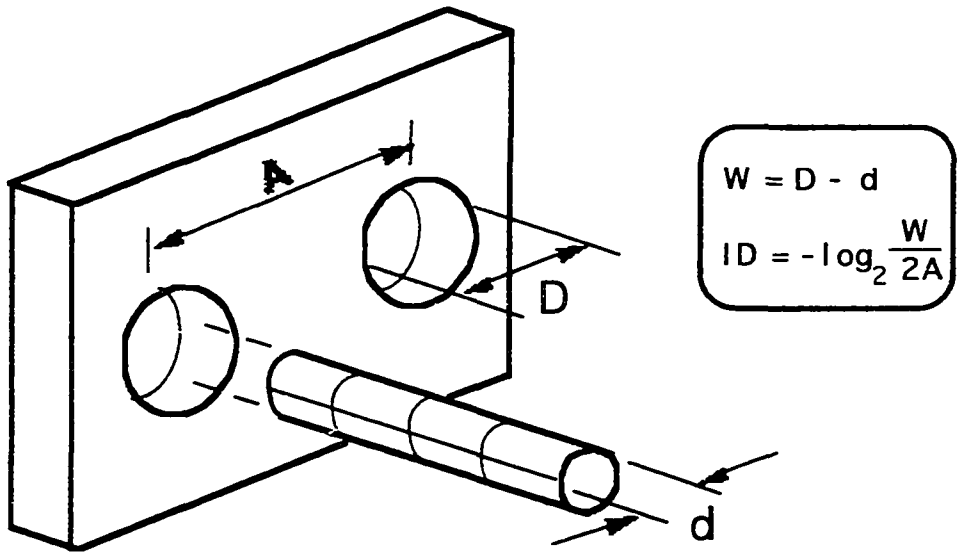

Figure 2.1: Measurements used to define task difficulty

Through an extensive review of human psychomotor, perceptual, and cognitive test batteries, the Naval Oceans Systems Center developed a peg insertion performance task specifically representative of teleoperator manipulative activities [Pepper, 1988; Spain, 1989]. The test battery requires subjects to move pegs of various diameters between holes of varied spacing. Movement times for peg motions are recorded and correlated with task difficulty. As defined by Fitts [1954], the binary Index of Difficulty, ID, for the one dimensional peg transfer task can be computed as:

$$ID = -log_2 \frac{W}{2A} \qquad \text{(bits/response)}$$

where A is the amplitude of the motion and W is the peg tolerance defined as the difference between the hole diameter and the peg diameter. These quantities are shown in Figure 2.1.

Fitts' Law [Fitts, 1954], relates task completion time to task ID by defining the movement time, mt, as follows:

$$mt = k_1 \, ID + k_2 \qquad \text{(sec)}$$

where k_1 and k_2 are characteristic constants of the individual operator and represent the slope and intercept of the Fitts' Law curve. The reciprocal of the slope of the Fitts' Law curve ($1/k_1$) has units (bits/sec) (which is identical to units of capacity for an information channel) and is thus an accepted measure of human information processing capacity [Remis, 1990]. Because the slope is a coarse measure and is susceptible to distortions resulting from changes in strategy between tasks of different difficulty, a more robust measure of processing capacity was also computed, called the *Binary Index of Performance*. This value, abbreviated as Ip, describes the information processing capacity required of the operator to perform a task of a particular difficulty. Ip also has units (bits/sec) and is defined as follows:

$$Ip = \frac{1}{mt} \log_2 \frac{W}{2A} \qquad \text{(bits/sec)}$$

or

$$Ip = \frac{ID}{mt} \qquad \text{(bits/sec)}$$

where mt is the movement time required to complete the task. Index of Performance Ip is an accepted measure of the information processing capacity of the teleoperator [Repperger and Remis, 1990].

2.2 EXPERIMENTAL HARDWARE

2.2.1 Virtual Fixture Performance Testbed

To implement a standardized Fitts' Law performance test in a telepresence environment, the following hardware subsystems were used: Peg Insertion Task Board, MBA Exoskeletal Master / Merlin Robot Arm Slave, Monocular Vision System, and Virtual Fixture Board. This equipment was assembled into a full system with facilities and support from Armstrong Laboratory, Human Sensory Feedback Group at Wright Patterson Air Force Base.

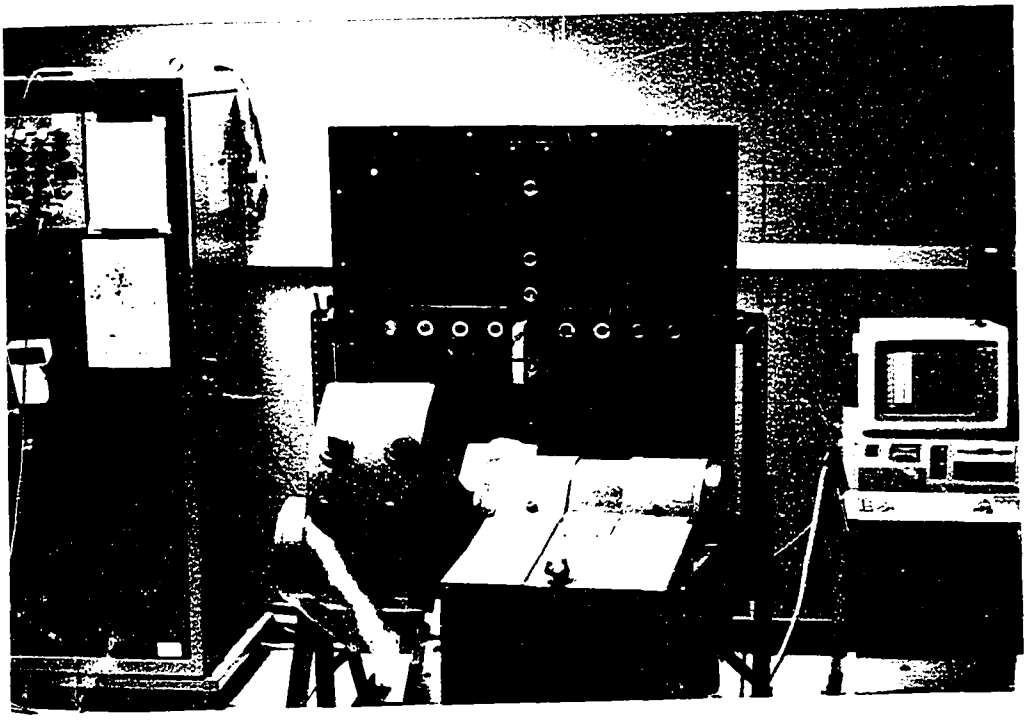

Figure 2.2: Photograph of the AF/Navy Teleoperator Performance Evaluation Battery Task Board used for implementing peg insertion Fitts' Law task. Also shown is the Merlin Robot Arm used as the slave in the teleoperated system.

Task Board: (AF/Navy Teleoperator Performance Evaluation Battery)

As shown in Figure 2.2, a specialized task board was used to implement the performance task. This task board was originally developed by the Naval Oceans Systems Center to implement a controlled peg-in-hole test battery for telemanipulation systems [Pepper, 1988; Spain, 1989]. Holes on the board are arranged to accommodate a wide range of peg sizes and movement amplitudes. This study made use of four holes (diameter 2.00 cm) spaced to allow horizontal peg movements of 4.00 cm and 16.00 cm. Three pegs (diameters 0.75 cm, 0.98 cm, 1.50 cm) were used in the task to vary the peg insertion tolerance. The bottom of each hole contains a microswitch connected to a PC via a Digital I/O controller card. This PC controls and monitors the peg insertion portion of the experiment. The computer is equipped with a real-time clock for measuring peg movement times to the nearest millisecond. The timer is started when the peg is extracted from the start hole and stopped when inserted into the target hole.

MBA Exoskeletal Master / Merlin Robot Arm Slave

As shown in Figure 2.3, the MBA Exoskeleton Master is a dual-arm, full upper body exoskeleton which can transduce motion in seven degrees of freedom for each arm [MB Associates, 1976]. In its current configuration, the device is used only as a sensor; no force information is reflected back to the user

through the exoskeleton. Thus, peg insertions were performed based only on visual and auditory feedback from the workspace. This reduced sensory feedback environment offered a good testbed for virtual fixturing.

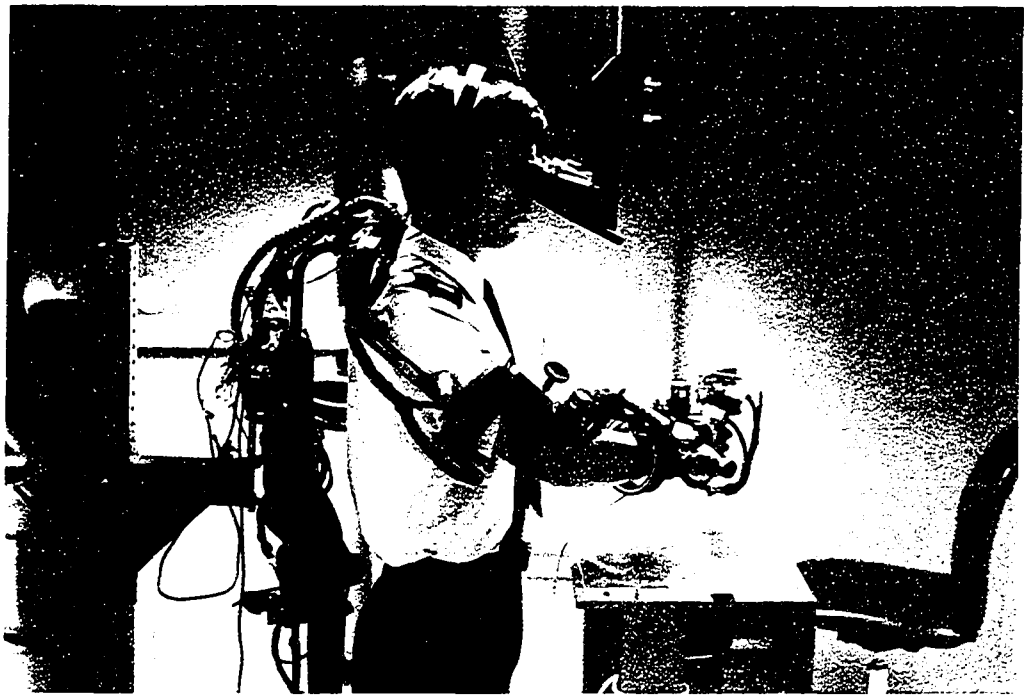

Figure 2.3: Photograph of subject wearing MBA kinesthetic exoskeleton used as the master in the teleoperator system.

In this experiment, the position of the subject's right hand was tracked and fed to the slave robot arm. An end-effector was constructed for the slave robot which housed the peg and allowed for easy interchange of the various diameter pegs. Because of the difficulty of the teleoperated peg insertion task without force feedback, the slave robot was programmed to maintain its end-effector orientation perpendicular to the board at all times. This arrangement simplified the task and eliminated the possibility of subjects torquing the peg in the hole. To further simplify the task, the robot was constrained to move only in a horizontal plane level with the target holes. Thus, the task only required two-dimensional positioning of the peg in the xy plane with x defined into the board, and y defined along the horizontal row of target holes as shown in Figure 2.4. It should be noted that although the task was constrained to planar positioning, it was sufficiently difficult to require subjects to practice for two 45-minute sessions before training was complete (i.e., until variation in trial performance dropped below 20%). The task was

designed to be simple in concept but difficult enough that the subjects would not approach the upper bound of the slave robot's ability to perform the task. In other words, the task was chosen so that performance bounds were a result of bandwidth limitations in the human operator, rather than in the physical hardware.

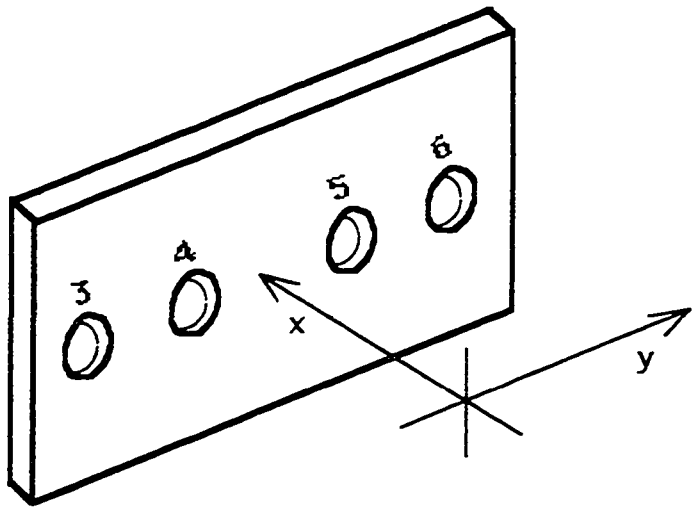

Figure 2.4: Schematic of Task Board with xy plane of allowable peg motion.

To protect the task board and the robot, the peg was rigidly constrained in all degrees of freedom except along the x axis pointing into the task board. Along that axis, the peg floated on a stiff spring so that strong forward impact of the peg with the board would not jar the board or damage the slave robot.

Monocular Vision System

A previous study using an identical task board performed by the Naval Oceans Systems Center has shown that stereo vision feedback offers no performance advantage over monocular projections used in this peg insertion task [Pepper, 1988; Spain, 1989]. Thus, a simple monocular vision system was implemented to provide visual feedback for this experiment. The system was designed as an inexpensive means of creating the illusion of operator presence while providing visual feedback of the workspace. The system used 7X power binocular optics focused upon a distant color video monitor which displayed the output of a single camera in the workspace focused on the task board. The monitor was placed at a distance from the optics in such a way that the magnification of the video image created the illusion that the task board was within reach of the operator's hands. Robot end-effector motion was scaled to match operator hand motion so that the apparent end-effector position corresponded to the user's kinesthetic sense of arm position. Thus when looking into the vision system, the operator saw the robot end effector in the location where his or her hand kinesthetically felt to be. Such a correlation between kinesthetic feedback and visual

feedback of robot end- effector position greatly enhanced the user's sense of presence within the workspace. Testing of the system revealed best results when the angle of the camera incident on the workspace was closely matched by the angle that the magnifying optics were incident upon the distant video monitor. When these angles were not similar, conflicting perspective cues hindered the illusion of presence. Figure 2.5 includes a rough schematic of the vision system as part of the overall system hardware.

Virtual Fixture Board

Rather than using a force reflecting exoskeleton to model the rigid impedance surfaces which compose the virtual fixtures, it was thought that the preliminary tests of virtual fixtures should not be influenced by the hardware limitations of current force reflecting devices. Thus, a Fixture Board was designed which allowed real rigid surfaces to be fabricated from acrylic sheets. The MBA exoskeleton was positioned to interact physically with the real acrylic surfaces and transfer the force information resulting from contact with those surfaces to the user. Because the user was only in physical contact with the exoskeleton and because the vision system prevented the user from seeing the acrylic surfaces which composed the fixture, the haptic sensations resulting from interaction with the surfaces were equivalent to what would have been generated by a high fidelity force reflecting exoskeleton. Thus, the haptic information presented to the user were just as "virtual" as if they were truly computer generated and reflected through the active elements of a force reflecting master. The benefit of this approach is that the reflected sensations of these surfaces was perfectly modeled. The surfaces felt crisp and real, free from the bandwidth and other limitations of most force reflecting devices. The drawback of this approach was that more abstract or complex fixture designs such as attractive or repulsive fields were impossible to generate using this method.

The virtual fixtures were constructed out of acrylic sheets and positioned upon a wooden platform in front of the user. To make the fixtures quickly interchangeable, they were constrained and positioned on the wooden platform by locating pins. The exoskeleton was fitted with a teflon cap at the end of the hand grip which was used as the contact surface between the exoskeleton and the fixture. The fixtures were treated regularly with an oil- based lubricant so that little friction was perceivable between the teflon and the acrylic surfaces. Although frictional surfaces could make effective virtual fixtures, the intent of this study was to look at very simple surfaces modeled only as rigid impedances free from any viscous or coulombic damping.

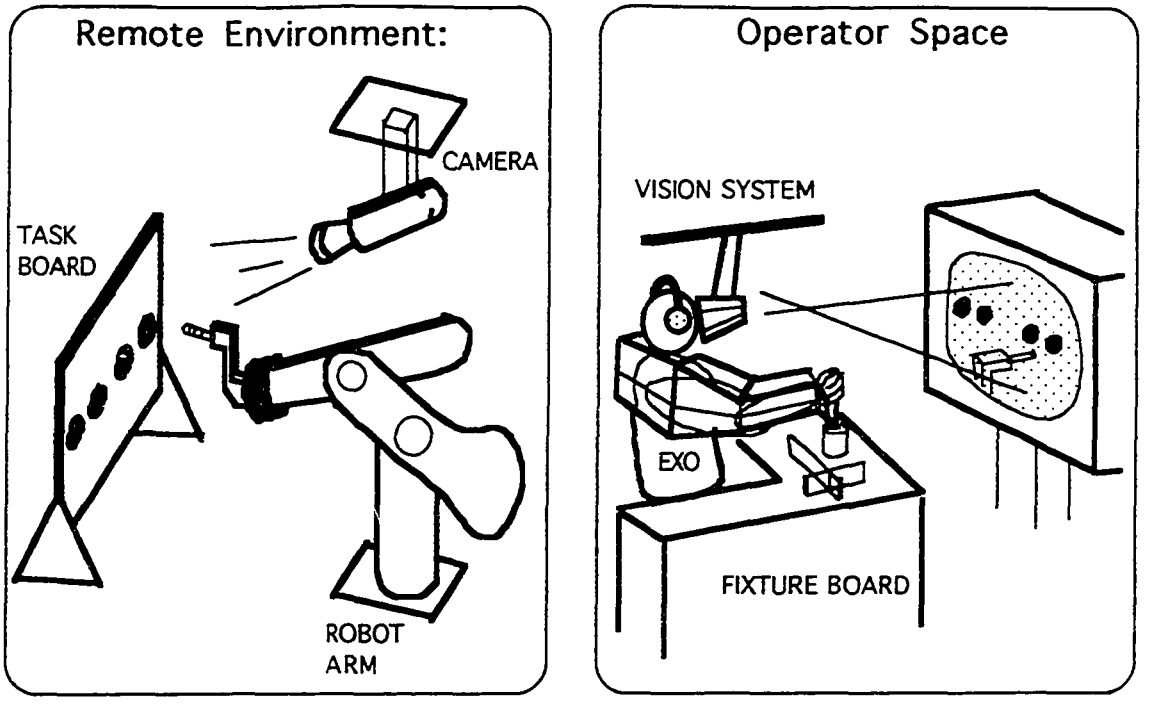

Figure 2.5: Complete Telepresence System developed to implement the testing of teleoperator performance in a standardized peg insertion task with and without the aid of virtual fixtures.

Although the haptic fixtures were modeled physically rather than computationally, the auditory fixtures tested were pure computer simulations. Simple compliant surfaces were modeled in which the compression of the surface was proportional to a linear change in pitch. Tones were generated on a PC and fed to the user via stereo headphones. These trials served as a testbed for alternate modality fixtures as well as a proof of concept of purely computer-generated fixtures.

2.2.2 Overall Test Setup

Having described each of the major components of this test setup, it is important to describe clearly the system as a whole before getting into the details of the subject testing. As shown in Figure 2.5, the system is divided into two physically separate parts: the *remote environment* and the *operator space*. The remote environment contains the task board, the Merlin robot arm, and a single video camera pointed at

task board. The camera is positioned so that the incident perspective is similar to what a human operator would see if standing directly in front of the board and performing the peg insertions in person. The operator space contains the exoskeleton, the vision system, and the virtual fixture board. Once inside the exoskeleton and vision system, the subject is presented with a projection of the image from the camera in the remote environment. The subject is given the illusion that the task board is situated directly before him, within reaching distance of the exoskeleton. In reality, the task board is on the opposite side of the laboratory, behind the subject and completely out of view. The fixture table is placed directly in front of the subject in such a way that it cannot be seen when looking through the vision system, but feels as though it occupies the same space as the apparent image of the task board. Thus, virtual fixtures implemented on the fixture board feel as though they are overlaid on top of their perception of the remote environment.

The remote robot arm is slaved to the right hand position of the exoskeleton. Thus, when the subject moves his right hand to interact with the image of the task board, the end-effector on the slave arm follows. Because the subject cannot see his own hand when looking into the vision system but does see the remote robot end-effector in the position where he feels his hand to be, a sense of presence within the remote environment is created. The subjects also wear a set of stereo headphones for use in implementing auditory perceptual overlays. Auditory surfaces and fields can be modeled on the control computer and interacted with by the user. The complete system provides a powerful testbed for projecting subjects into a remote environment and overlaying haptic and auditory information on top of the reflected percepts.

23 EXPERIMENT DESIGN & PROCEDURE

2.3.1 Experiment Design

The goal of this study was to investigate the effect that the presentation of perceptual overlays in the form of virtual fixtures has upon performance in a standard telemanipulation task. Eight simple fixtures were developed for comparison with a no-fixture control case. Six of the test fixtures were purely haptic constructs while two fixtures introduced both haptic and auditory sensory information. It should be noted that the design of these test fixtures was *not* motivated by finding the most effective perceptual aid for this particular peg insertion task. This particular set of test fixtures was chosen to provide general insight into the design and use of virtual fixtures and to allow the evaluation of some of the *basic building blocks* from which fixtures can be made.

When designing the test fixtures, the peg insertion task was not thought of as a single motion, but rather as a combination of two phases: ballistic motion and fine positioning. The ballistic phase was defined as the time from which the peg is removed from the start hole until it makes first contact with the target hole. The fine positioning phase was defined to begin at the end of the ballistic motion and continue until the peg is properly positioned in the hole [Remis and Repperger, 1990]. This decomposition of the peg insertion task was found to be a useful conceptual guide for fixture design. Some of the test fixtures were targeted at aiding ballistic motion, some were targeted at enhancing fine motion, and some were intended to help both. Figure 2.6 shows a schematic representation of each of the eight fixtures tested. These fixtures are shown overlaid on top of the task board as they are perceived by the subjects. Although these fixtures are represented graphically in this figure, they are perceived only as forces or sounds by the subjects. Although instilling visual qualities to fixtures is a viable application of perceptual overlays and an important topic for future investigation into virtual fixtures, this study was restricted only to haptic and auditory sensory modalities.

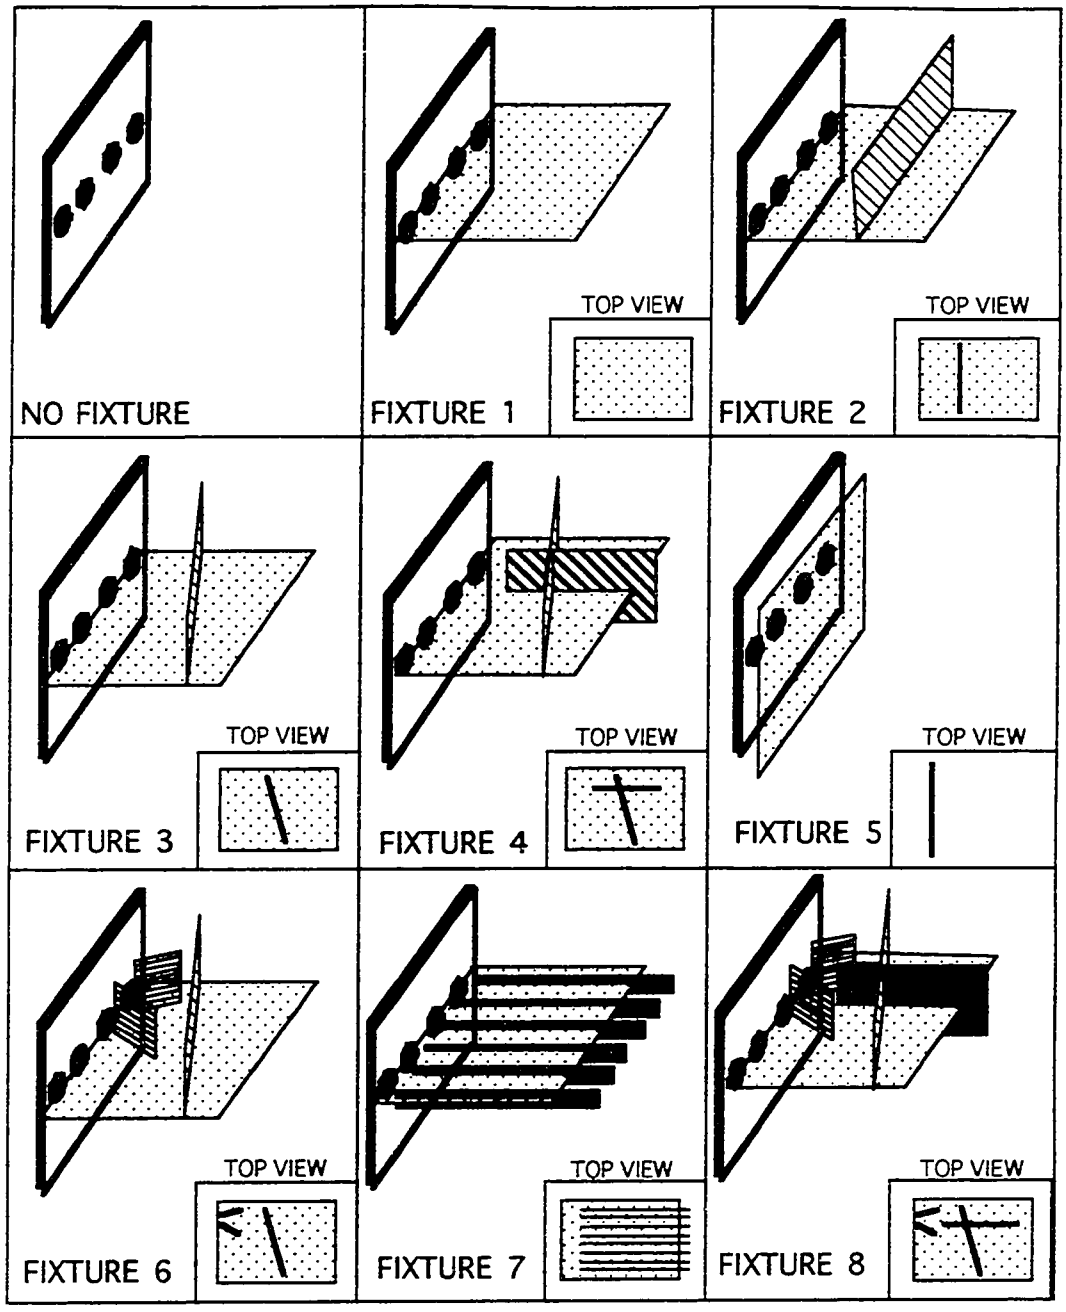

Figure 2.6: Virtual fixtures shown projected on top of the task board as perceived by the subjects. Also shown are top views, looking down from above the task board. The fixtures are shown graphically here, but they are perceived haptically or auditorily by subjects.

As shown in Figure 2.6, the eight fixtures tested are composed of simple combinations of planar surfaces. Fixture 1 is simply a rigid horizontal surface oriented like a table top in the workspace and positioned so that contact with the surface will result in vertical alignment of the operator's hand with the holes in the peg board. Fixture 2 is just like Fixture 1, but includes a second surface parallel to the plane of the task board and located three inches back from the board. When using this fixture, subjects operate in the space between the vertical plane and the task board. This arrangement provides some limits upon the operator's ballistic motion. Fixture 3 is similar to Fixture 2, except the vertical plane is angled so as to converge with the task board as the target hole is approached. This arrangement acts like a funnel, confining and guiding ballistic motion towards the target hole.

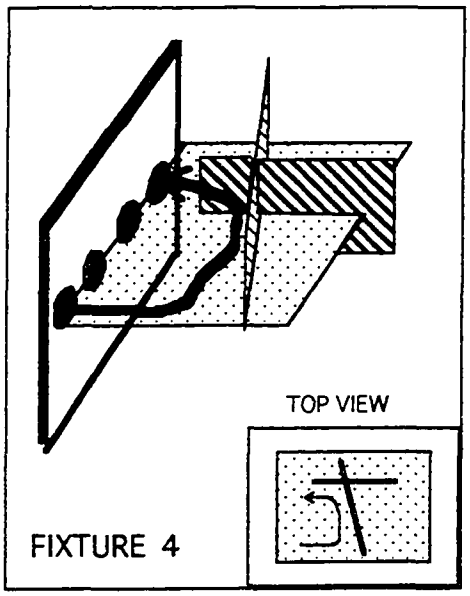

Figure 2.7: Sample trajectory of peg from start hole to target hole as guided by a virtual fixture. As shown, the fixture does not define the motion of the operator's hand, but rather defines the boundaries of the motion.

Fixture 4 is identical to Fixture 3 except a third plane is added directly in front of the target hole. This fixture influences the ballistic phase even further, stopping hand motion when close to the target hole. Fixture 6 is similar to Fixture 3, except a second funnel-like fixture is added in front of the target hole. Thus, ballistic motion is guided as was done by Fixture 3, and fine motion is guided by the additional surfaces. Figure 2.7 shows a sample trajectory superimposed upon Fixture 4, making the fixture implementation clear. This drawing demonstrates how the use of a fixture might influence the trajectory of the user's hand. As shown in the figure, the rigid surfaces do not define the motion of the peg, but rather influence the operator's trajectory by confining the boundaries of peg motion.

Fixture 5 is very different from those presented thus far in both its geometry and implementation. The key is that this fixture was *not* interacted with by the subject's right hand (the hand that performed the peg insertion task), but rather was designed for interaction only with the *unused* left hand. Fixture 5 is a rigid impedance plane parallel to the task board and located approximately 0.5 inches in front of the board. The subject would place his left hand upon the planar surface while performing the task with the right hand. The intent of this fixture was to isolate the effect of localization on performance. It was hypothesized that all fixtures provided some localizing information to the users which enhanced their understanding of the geometry of the workspace and allowed them to better correlate their kinesthetic sense of hand position to the remote site. Because interaction with other fixtures would guide or limit the operator's motion, it was impossible to isolate performance increases due to the localizing effect alone. Thus, Fixture 5 was designed to interact only with the unused hand and therefore only effects performance by providing localizing cues.

Fixture 7 and Fixture 8 are the only fixtures which implement both haptic and auditory information. Fixture 7 is identical to Fixture 1, but also introduces a texture-like field of auditory information. The information is represented as a series of surfaces perpendicular to the task board as shown in black in Figure 2.6. Operator interaction with these surfaces results in the production of an audible tone. The pitch of the tone increases from left to right across the task board. Fixture 8 is identical to Fixture 6, but it introduces a single compliant surface in front of the target hole as shown in black in Figure 2.6. This surface is modeled as a compliant surface with a proportional stiffness such that interaction with this surface produces a tone proportional to compression.

2.3.2 Experimental Protocol

A series of tests was run to evaluate subject performance using each fixture configuration. Operator performance was recorded during test periods which included 12 practice and 36 timed peg insertion trials for each fixture studied. A single trial consisted of moving a peg from a designated starting hole to a designated target hole. The holes were referenced by numerals (3, 4, 5 or 6) located above each hole as shown in Figure 2.4. Two different peg motions were studied in these tests: a 16 cm motion from hole 3 to hole 6 and a 4 cm motion from hole 5 to hole 6. The 36 trial period was divided into three groups of 12 trials. Each of these groups required the subject to perform the insertion task using a different peg size. After the completion of each group of 12 trials, subjects were instructed to rest for approximately two minutes while a new peg size was installed in the robot end-effector.

The use of two motion amplitudes (4 cm and 16 cm) and three peg sizes (0.75 cm, 0.98 cm, 1.50 cm diameter pegs) allowed for the testing of insertion trials with six different task difficulties. Table 2.1 shows all combinations of peg size and motion amplitude and lists the Index of Difficulty for each task as dictated by Fitts' Law [Fitts, 1954]. The order in which the three peg sizes were presented to each subject was randomized to ensure that mental and physical fatigue had similar effects on all task difficulties.

Table 2.1: *Index of Difficulties* (ID) shown for all combinations of peg size and peg motion amplitudes as predicted by Fitts' Law.

	Motion Amplitude	
Diameter	5 to 6 (4 cm)	3 to 6 (16cm)
0.75 cm	**2.68** bits/response	**4.68** bits/response
0.98 cm	**2.98** bits/response	**4.98** bits/response
1.50 cm	**4.00** bits/response	**6.00** bits/response

Subjects were instructed to begin each test period with a hole 3 to hole 6 peg insertion trial. Subjects would then perform a hole 5 to hole 6 peg insertion trial, then a 3 to 6 peg insertion trial, and cycle in that manner throughout the 36 trials. Subjects were allowed to proceed through the test period at their own pace. To automate the testing procedure and allow subjects to proceed through the test with little operator intervention, a number of simple beeps was used to guide subject activity. When a subject inserted a peg into the correct starting hole, the task board control computer would emit an audible tone to signal that the peg was properly positioned. The subject would then keep the peg in the start hole for two seconds until the control computer produced a second tone which signaled that the task could now be performed at will. The reason for this two second waiting period was to ensure that the subject maintained a steady contact with the microswitch at the base of the start hole. The subjects were *not* required to perform the insertion task as soon as the second tone was heard; this second tone simply meant "you are free to go whenever ready." Thus the subjects could mentally prepare themselves for the particular insertion task and begin at will. When ready, the subject would remove the peg from the start hole and insert it into the target hole as fast as possible. Upon insertion into the target hole, the control computer would emit a tone to signal that the trial had been successfully completed. The subject was then free to proceed to the next insertion trial at will. At the end of each trial the control computer recorded the movement time along with the peg size and motion amplitude. After the completion of each group of 12 trials, the control computer would emit a long-duration tone which signaled the subject to rest for two minutes while the

operator exchanged peg sizes. Post testing interviews revealed that all subjects were comfortable with the use of audible tones to automate the testing procedure.

Subjects were tested over 9 experimental sessions, each lasting 45 to 60 minutes. To minimize the effects of fatigue and boredom, no subject completed more than one experimental session during a single day of testing. Each of the first two experimental sessions included two test periods of 36 trials. The first period of each session was performed with no fixture and the second was performed with Fixture 1. These initial 144 trials were treated only as practice during which the subjects familiarized themselves with the use of the exoskeleton, merlin robot arm, vision system, and fixture table. It was found that by the end of the second practice session, all subjects had sufficiently learned the task that variability in movement times for trials of the same difficulty had fallen below 20% for every subject, with a mean variability of 14% for all subjects. Once learning had stabilized, subjects were sequentially tested using each of the test fixtures. One or two new fixtures were tested during each daily session in addition to a baseline fixture which was tested during every session. Fixture 1 was the baseline fixture and was used to track performance increases resulting from day-to-day learning so that any such effects could be compensated. Table 2.2 lists all of the daily sessions along with the fixtures tested during that session.

Table 2.2: *Testing Schedule:* All seven tests shown along with corresponding daily sessions and fixtures tested on that day.

TEST	SESSION	FIXTURES STUDIED
Practice	day 1	No Fixture / Fixture 1
Practice	day 2	No Fixture / Fixture 1
Test I	day 3	No Fixture / Fixture 1
Test II	day 4	Fixture 1 / Fixture 2 / Fixture 3
Test III	day 5	Fixture 1 / Fixture 4
Test IV	day 6	Fixture 0 / Fixture 5
Test V	day 7	Fixture 1 / Fixture 6
Test VI	day 8	Fixture 1 / Fixture 8
Test VII	day 9	Fixture 1 / Fixture 7

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

2.4 RESULTS

The following section presents all results of fixture performance comparisons in tabular form. The basic quantities presented are the recorded movement times and the coefficients of variation of movement times for trials of the same difficulty. The coefficient of variation is a measure of subject consistency at a particular task and is an indication of the repeatability of the results. Also computed is the slope of a linear regression line relating movement time to index of difficulty. The inverse of the slope has units (bits/sec) and is a measure of the operator information processing capacity. Because the slope is susceptible to distortions resulting from changes in strategy between tasks of different difficulty, a more robust measure of processing capacity, called the *Binary Index of Performance*, was also computed . This value, abbreviated as Ip, describes the information processing capacity required of the operator to perform a task.

Tables 2.3 through 2.5 list the mean movement time results for each fixture tested. Tables 2.6 through 2.8 list the computed Index of Performance (Ip) for each fixture tested. Table 2.9 lists the coefficients of variation, slopes, and the inverse of the slopes for each fixture tested. Figure 2.8 graphically compares the performance, showing the mean movement time versus index of difficulty curves for all test fixtures.

TABLE 2.3: *Mean Movement Times* for each Index of Difficulty are shown for eight subjects using no fixture as well as using Fixtures 1, 2, 3, 4 and 6. Also computed is the mean movement time across all trials (across task difficulty).

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
No Fixture	1123ms	1280ms	1353ms	1946ms	1866ms	2162ms	1622ms
Fixture 1	957ms	977ms	1118ms	1595ms	1656ms	1758ms	1344ms
Fixture 2	842ms	883ms	1030ms	1435ms	1527ms	1631ms	1225ms
Fixture 3	682ms	704ms	795ms	1293ms	1350ms	1469ms	1049ms
Fixture 4	692ms	713ms	815ms	1376ms	1471ms	1543ms	1102ms
Fixture 6	667ms	722ms	835ms	1301ms	1385ms	1495ms	1068ms

TABLE 2.4: *Mean Movement times* for each Index of Difficulty are shown for eight subjects using no fixture as well as using the localizing fixture 5 with the unused hand. Also computed is the mean movement time across all trials (across task difficulty).

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
No Fixture	833ms	927ms	991ms	1617ms	1600ms	1666ms	1272ms
Fixture 5	656ms	662ms	807ms	1448ms	1396ms	1497ms	1078ms

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 2.5: *Mean Movement Times* for each Index of Difficulty are shown for eight subjects using no fixture, using purely haptic Fixtures 1 and 6, as well as using auditory/haptic Fixtures 7 and 8. Also shown is the mean movement time across all trials. It should be noted that Fixtures 1 and 7 are identical except for the addition of auditory information in 7. The same is true for Fixture 6 and Fixture 8.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
No Fixture	1073ms	1177ms	1259ms	1845ms	1832ms	2058ms	1541ms
Fixture 1	840ms	910ms	969ms	1539ms	1569ms	1674ms	1250ms
Fixture 7	687ms	727ms	718ms	1310ms	1325ms	1369ms	1023ms
Fixture 6	641ms	691ms	811ms	1266ms	1361ms	1461ms	1039ms
Fixture 8	614ms	610ms	701ms	1136ms	1154ms	1270ms	915ms

TABLE 2.6: *Index of Performance*, Ip = (ID / mean movement time), is computed and shown for 8 subjects using no fixture as well as using Fixtures 1, 2, 3, 4 and 6.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean Ip
No Fixture	2.34	2.33	2.96	2.40	2.67	2.77	2.58
Fixture 1	2.80	2.99	3.58	2.93	3.01	3.41	3.12
Fixture 2	3.18	3.37	3.88	3.26	3.26	3.68	3.44
Fixture 3	3.93	4.23	5.03	3.62	3.69	4.08	4.10
Fixture 4	3.87	4.18	4.91	3.40	3.39	3.89	3.94
Fixture 6	4.02	4.13	4.79	3.60	3.60	4.01	4.02

TABLE 2.7: *Index of Performance*, Ip = (ID / mean movement time), is computed and shown for eight subjects using no fixture as well as using the localizing Fixture 5 with the unused hand.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean Ip
No Fixture	3.22	3.21	4.03	2.89	3.11	3.60	3.60
Fixture 5	4.09	4.50	4.96	3.23	3.58	4.01	4.01

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 2.8: *Index of Performance.* Ip = (ID / mean movement time) , is computed and shown for eight subjects using no fixture, using purely haptic Fixtures 1 and 6, as well as using auditory/haptic Fixtures 7 and 8. Ip has units (bits/sec) and is an accepted measure of human information processing capacity. It should be noted that Fixtures 1 and 7 are identical except for the addition of auditory information in 7. The same is true for Fixture 6 and Fixture 8.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean Ip
No Fixture	2.50	2.67	3.18	2.54	2.72	2.92	2.76
Fixture 1	3.19	3.27	4.12	3.04	3.17	3.58	3.40
Fixture 7	3.90	4.01	5.57	3.57	3.76	4.38	4.20
Fixture 6	4.18	4.31	4.93	3.70	3.66	4.11	4.15
Fixture 8	4.36	4.88	5.71	4.12	4.32	4.72	4.69

Table 2.9: *Mean Coefficient of Variation* for all subjects performing tasks of the same difficulty with each fixture is shown. Also shown is the slope and standard error of a linear regression analysis of movement time against index of difficulty. The reciprocal of the slope is also shown and has units (bits/sec). This value is an accepted measure of information processing capacity.

FIXTURE	MEAN C.O.V.	SLOPE [ms/bit]	Std Err of Coef	1/SLOPE [bits / sec]
No Fixture	13.7% ±2.5	317	43	3.15
Fixture 1	11.9% ±2.0	276	39	3.62
Fixture 2	12.8% ±3.2	267	37	3.75
Fixture 3	11.1% ±2.7	271	47	3.69
Fixture 4	8.6% ±1.9	300	56	3.33
Fixture 5	13.0% ±1.3	302	62	3.31
Fixture 6	9.6% ±2.8	280	44	3.57
Fixture 7	11.2% ±1.9	244	61	4.10
Fixture 8	9.2% ±1.5	228	40	4.39

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

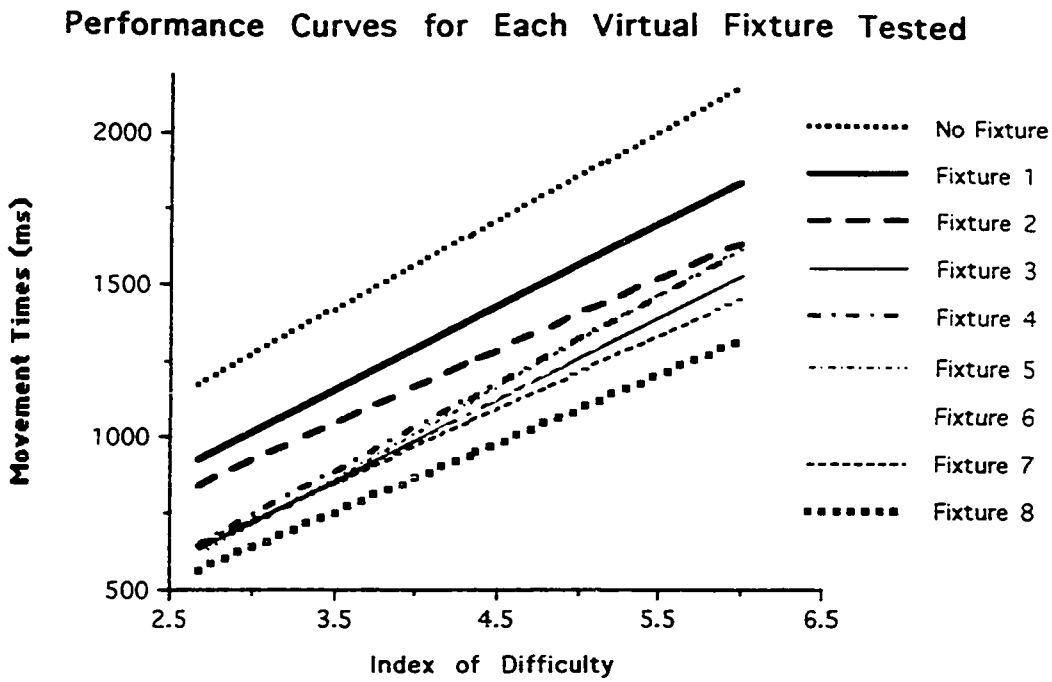

Figure 2.8: Performance using each fixture shown for comparison as best line fit to Fitts' Law relation between mt and ID

In order to quantify performance increase due to fixture use in a meaningful way that is resistant to variations across task difficulty, percentage changes in movement times and Index of Performance were computed for each fixture with respect to the no-fixture case. The percentage decrease in movement times is a unitless measure of the effectiveness of a fixture to speed operator performance in a task. The percentage increase in Index of Performance is a unitless measure of the increase in the information processing capacity of the operator or the decrease in processing requirements of the task resulting from fixture use. Tables 2.10 and 2.11 list the percentage decrease in movement times corresponding to the use of each fixture. Tables 2.12 and 2.13 list the percentage increase in processing capacity Ip associated with the use of each fixture.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 2.10: *Percent Decrease in Movement Time* for each fixture with respect to the no-fixture case. Values are computed for each index of difficulty. Also shown is the mean percentage decrease in movement time for each fixture. This mean value represents the effectiveness of a fixture to enhance operator performance.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
Fixture 1	15.96%	26.85%	19.02%	19.82%	11.93%	20.61%	19.03% ±4.9
Fixture 2	28.60%	36.71%	27.11%	30.23%	19.98%	28.00%	28.44% ±5.4
Fixture 3	48.86%	58.06%	51.96%	40.32%	32.09%	38.17%	44.91% ±9.1
Fixture 4	47.49%	56.90%	49.63%	34.32%	23.67%	33.41%	40.90% ±11.9
Fixture 5	23.77%	33.35%	20.47%	11.03%	13.62%	10.69%	18.82% ±7.9
Fixture 6	50.95%	55.74%	47.35%	39.73%	29.59%	36.48%	43.31% ±9.2
Fixture 7	43.86%	47.27%	54.73%	33.91%	32.12%	40.21%	42.02% ±8.4
Fixture 8	54.42%	63.46%	56.94%	47.57%	45.41%	47.36%	52.52% ±6.5

TABLE 2.11: *Percentage Decreases in Movement Times* when using Fixtures 7 and 8 as compared to Fixtures 1 and 6 respectively. Since Fixtures 7 and 8 are identical to 1 and 6 in all ways except for the addition of auditory information, the percentage decrease in movement times reflects the performance advantage associated with overlaying auditory information.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
Fixture 7	20.04%	22.36%	29.76%	16.08%	16.86%	20.05%	20.89% ±4.9
Fixture 8	4.30%	12.45%	14.55%	10.82%	16.46%	13.99%	12.10% ±4.2

TABLE 2.12: *Percent Increase in Processing Capacity* Ip for each fixture with respect to the no-fixture case. Values are computed for each index of difficulty. Also computed is the mean percentage increase in processing capacity associated with each fixture.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
Fixture 1	19.66%	28.33%	20.95%	22.08%	12.73%	23.10%	21.14% ± 5.0
Fixture 2	35.90%	44.64%	31.08%	35.83%	22.10%	32.85%	33.73% ± 7.4
Fixture 3	67.95%	81.55%	69.93%	50.83%	38.20%	47.29%	59.29% 15.3
Fixture 4	65.38%	79.40%	65.88%	41.67%	26.97%	40.43%	53.29% ±18.9
Fixture 5	27.02%	40.19%	23.08%	11.76%	15.11%	11.39%	20.31% ±10.8
Fixture 6	71.79%	77.25%	61.82%	50.00%	34.83%	44.77%	56.74% ±15.3
Fixture 7	56.00%	50.19%	75.16%	40.55%	38.24%	50.00%	51.69% ±13.2
Fixture 8	74.40%	82.77%	79.56%	62.20%	58.82%	61.64%	69.90% ± 9.5

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 2.13: *Percentage Increases in Processing Capacity* Ip when using Fixtures 7 and 8 as compared to Fixtures 1 and 6 respectively. Since Fixtures 7 and 8 are identical to 1 and 6 in all ways except for the addition of auditory information, the percentage increase in Ip reflects the performance advantage associated with overlaying auditory information.

	INDEX OF DIFFICULTY						
FIXTURE	2.68	2.98	4.00	4.68	4.98	6.00	mean
Fixture 7	22.26%	22.63%	35.19%	17.43%	18.61%	22.35%	23.08% ±6.3
Fixture 8	4.31%	13.23%	15.82%	11.35%	18.03%	14.84%	12.93% ±4.7

2.5 DISCUSSION

The following section sequentially addresses the change in performance associated with the use of each virtual fixture design implemented in this set of experiments and discusses the implications. The objective of this analysis is to reveal empirically if perceptual overlays can enhance performance in telepresence tasks. It is hoped that this discussion with reveal general insights into how these various embellishments of a user's perceptual reality can effect operator performance and understanding within a remote environment.

Before discussing the performance results associated with the use of each virtual fixture configuration, the issue of repeatability and reliability of the results should be addressed. As shown in Table 9, the mean coefficients of variations for task completion is under 14% for all virtual fixture configurations, with a mean variability of 10%. This implies that for any given peg insertion task, of a particular index of difficulty, performed by a given subject, the variation in performance for all such trials was on average 10%. Such consistency for human performance is surprisingly good and suggests that the results of these tests are highly reliable. Such consistency further suggests that all subjects were adequately trained during the practice session portion of the testing procedure and that learning had little effect on performance variations during the data collection trials.

2.5.1 Fixture 1

As shown in Figure 2.6, virtual fixture 1 is a rigid surface oriented perpendicularly to the plane of the task board. When interacting with this fixture, the subject's hand motion in restrained to move only in

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

the plane of the peg holes. Thus, this fixture reduces the Cartesian degrees of freedom on hand motion from three to two. An interesting point about this fixture is that the end-effector of the slave robot is restricted by software to planar two degree of freedom motion regardless of the operator's commands. Thus, the degree of freedom which is eliminated by the use of fixture 1 has no effect on the commanded position of the slave robot arm. Since this degree of freedom plays no part in task execution, it would seem that fixture 1 would have little effect on the performance of the operator.

Comparing Indexes of Performance (Ip) as shown in Table 2.6 and Table 2.12, operator performance when using virtual fixture 1 (measured in terms of the human information processing capacity) was 21% (±5%) greater than with no fixture. Thus, these results demonstrate that a perceptual overlay as simple as a flat surface haptic percept can enhance operator performance in a telemanipulation task. Deeper comparison between the fixture 1 trials and the no fixture trials reveals an interesting insight. In both cases, the remote peg insertion task was a 2 degree of freedom manual operation. Without using a fixture, the subjects were free to move their hand in all three Cartesian directions even though they were well aware that the slave robot's position was locked in the plane of the task. When using virtual fixture 1, the subjects were guided to stay in the plane. This implies that even after the 144 trials of practice, the subjects were unable to completely ignore the irrelevant, un-tracked degree of freedom and were wasting mental processing capacity on it. In other words, even through the task was simplified by eliminating one degree of freedom from the human-manipulator mapping through a simple software fix, the addition of a perceptual overlay to provide the operator with a perceptual representation of that software fix was effective in boosting performance by more than 20%. This is a strong confirmation that perceptual overlays do more than simplify task requirements, they provide an embellished perceptual reality which represents a given task simplification in a form that humans can readily process.

Consider again the telepresence surgey example proposed in Chapter 1 which described a hypothetical virtual fixture composed of a flat plane passing through a patients body to prevent incisions below a particular depth. The results of testing with virtual fixture 1 suggests that such a fixture would not only limit motion orthogonal to the plane, it would also enhance operator performance parallel the plane. Thus, not only would the surgeon be safe from encountering vital organs located beneath the protective surface, the surgeon would have enhanced dexterity above the surface.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

2.5.2 Fixture 2

As shown in Figure 2.6, virtual fixture 2 is the same as fixture 1 except a second rigid surface is added parallel to the surface of the task board. Like fixture 1, fixture 2 restricts operator hand motion to the horizontal plane. The added surface also restricts how far back from the board the operator can draw the peg. Comparing Indexes of Performance (Ip) as shown in Table 2.12, operator performance when using virtual Fixture 2 was 34% greater than in the no fixture trials. This comparison clearly shows that the use of fixture 2 significantly enhanced performance over the no fixture case and that use of the additional surface in fixture 2 as compared to fixture 1 further heightened performance improvement.

The question thus remains, why did the additional surface in fixture 2 increase operator performance? It is believed that the additional surface served a number of purposes by restricting how far back the peg could be drawn from the board. First, this surface restricts operator ballistic motion, preventing unnecessarily wide trajectories. This should simplify the operators conceptualization of the task (i.e. reduce the information processing demands of the task) because distance from the task board is no longer a parameter the operator needs to be constantly concerned with. Secondly, this surface acts as a localizing agent, providing haptic cues which give the user a better sense of hand position in the remote workspace. In perceptual terms, fixture 2 introduces a haptic percept into the operator's perceptual reality which may enhance operator presence in the remote site by providing an additional modality through which efference(hand motion) can be related to afference (haptic feedback) during task performance. Finally, by providing depth information haptically, the use of fixture 2 reduces the demand on the visual system to gage depth and frees up that modality for other uses such as tracking the target hole.

2.5.3 Fixture 3

As shown in Figure 2.6, virtual fixture 3 is the same as fixture 2 except the second rigid surface is not parallel to the surface of the task board, but rather is diagonal. Like fixture 1, fixture 3 maintains operator hand motion in the plane of the insertion task. Like Fixture 2, the additional surface restricts how far back from the board the operator can draw the peg. The uniqueness of fixture 3 is that it does not uniformly restrict how far back the peg can be drawn, but rather converges hand motion toward the board as the target hole is approached. Comparing Indexes of Performance (Ip) as shown in Table 2.6 and Table 2.12, operator performance when using virtual Fixture 3 was 59% greater than with no fixture. Like fixture 2, fixture 3 introduces haptic overlays into the operator's perceptual reality which may enhance localization to the remote workspace as well as improve the sensation of operator presence. The primary advantage of this fixture is that it guides ballistic motion more efficiently, funnelling the gross motion towards the

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

target hole. It is believed such assistance in target convergence reduces the demand upon kinesthetic and visual feedback and thus reduces the information processing requirements of the task.

2.5.4 Fixture 4

As shown in Figure 2.6, virtual fixture 4 is the same as fixture 3 except for the addition of a third rigid surface which crosses the diagonal surface. This additional surface was placed such that the operator would contact this surface when the peg was directly in front of the target hole. Whereas fixture 3 aided ballistic motion by converging hand position towards the target hole, fixture 4 was intended to provide further trajectory shaping by halting hand motion when the peg was aligned with the hole. Comparing Indexes of Performance (Ip) as shown in Table 2.12, operator performance when using virtual fixture 4 was 53% greater than in the fixture trials. The results for fixture 4 are not significantly different from those for the fixture 3 case, suggesting that the addition of the third surface did not enhance performance in these trials. Thus we can conclude that the additional haptic percept added to the operator's perceptual reality was not useful for this particular task. An interesting side note is that in a time delay study (to be discussed in Chapter 3) fixture 4 was found to provide significant improvement over fixture 3 when communication delays were introduced into the system. Thus, under certain task conditions, this particular enhancement of the user's perceptual relality is useful in performance enhancement. This makes sense because additional haptic references become more important when time delays distort the relation between kinesthetic sence of hand position and visual feedback of manipulator position.

2.5.5 Fixture 5

As shown in Figure 2.6, virtual fixture 5 is composed of a single rigid surface which is parallel to the task board and is located approximately 0.5" in front of the board. Whereas all fixtures thus far described were designed for interaction with the operator's right hand (the hand that manipulates the peg), fixture 5 was designed for interaction only with the operator's unused hand. Subjects place the palm of their left hand upon the virtual surface while they perform the task with their right hand. The purpose of testing this fixture was to isolate the effect of localization upon performance because this fixture can only influence performance through pure localization. The idea for this fixture was developed while observing operators perform the peg insertion task in person (not telepresent). All subjects rested their unused left hand upon the task board while performing the task with their right hand. The hypothesis was that operators gather useful perceptual information through the unused hand which may act to localizes the user with respect to the worksite. Comparing Indexes of Performance (Ip) as shown in Table 2.12, operator performance when using virtual fixture 5 was 20% greater than in no-fixture trials. These results suggest that kinesthetic/haptic localization to the remote site plays a part in the effectiveness of virtual fixtures to

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

enhance performance. These results further suggest the importance of telepresence systems which allow for dual arm interaction with the workspace even if for tasks which only require one hand.

2.5.6 Fixture 6

As shown in Figure 2.6, virtual fixture 6 is identical to fixture 3 with the addition of two angled surfaces which guide hand motion directly into the target hole. Thus fixture 6 guides ballistic motion by converging gross hand trajectory towards the task board as well as guiding fine motion by funneling hand position into allignment with the target. Although this fixture is composed of 4 intersecting planes, it can be thought of as two cones projected into the two dimensional workspace. One cone funnels ballistic motion across the length of the board; the other cone funnels fine motion into the target hole. Comparing Indexes of Performance (Ip) as shown in Table 2.12, the operator performance when using virtual fixture 6 was 57% greater than in the no-fixture trials. The results for fixture 6 are well within standard deviation of those for the fixture 3 case, thus suggesting that the addition of the fine positioning surfaces did not enhance performance in these trials. Thus it seems that the additional haptic cues which fixture 6 introduced into the user's perceptual reality were not helpful for the given task. An interesting side note is that in a time delay study (presented in Chapter 3), fixture 6 was found to provide significant improvement over fixture 3 when communication delays were introduced into the system. This makes sense because target aquision becomes more difficult when time delays are present.

2.5.7 Fixture 7

As shown in Figure 2.6, virtual fixture 7 is identical to fixture 1 with the addition of field-like auditory feedback into the operator's perceptual reality. The sonic information is modeled as a series of parallel planes such that crossing of each plane results in a tone of a given frequency. As planes are crossed from left to right, the frequency associated with each plane increases. Both the density of auditory pulses and the pitch of each pulse are useful cues. The result is a texture-like sonic overlay which provides rich position, velocity, and acceleration feedback to the operator. Comparing Indexes of Performance (Ip) as shown in Tables 2.12 and 2.13, operator performance when using virtual fixture 7 was 23% greater than the fixture 1 case and 52% greater than with no fixture. The over 20% improvement in performance when using fixture 7 over fixture 1 (which differ only by the addition of auditory information in 7) suggests that the use of multiple sensory is thus a powerful tool in fixture design. It further suggests that overlaying an auditory field on top of the workspace to provide an abstract perceptual representation of hand position, velocity, and acceleration is a simple means of enhancing operator performance by embellishing a user's perceptual reality. Post testing subject interviews revealed that all reported having a better sense of the shape of their peg insertion trajectories as a result of this perceptual overlay.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

2.5.8 Fixture 8

As shown in Figure 2.6, fixture 8 is identical to fixture 6 with the addition of an auditory compliant surface modeled such that compression of the surface corresponds to increasing pitch of the auditory feedback. The surface is positioned such that first contact occurs just before the peg should be inserted into the target hole. This auditory surface provides a well positioned localizing reference to act as a perceptual cue to indicate that the peg has traveled far enough and now should be inserted. Comparing Indexes of Performance (Ip) as shown in Tables 2.12 and 2.13, operator performance when using virtual fixture 8 was 13% greater than the fixture 6 case and 70% greater than with no-fixture. The improvement in performance when using fixture 8 over fixture 6 (which differ only by the addition of auditory information in 8) suggests that the use of the additional modality was useful in increasing information processing capacity. It should also be noted that of all the fixtures tested, the use of fixture 8 resulted in the greatest improvement in movement times and Index of Performance over the no-fixture case. This suggests that the use of multiple modalities in fixture design is a powerful perceptual aid and should be expanded to visual and tactile information as well. Post testing interviews revealed that the use of the auditory fixtures caused subjects to alter their conceptualization of the task such that a successful trial no longer just looked and felt a certain way, but also sounded a certain way.

We can further compare the advantage of using multiple modalities in fixture design by analyzing the slopes of the Fitts' Law curve relating movement times to Index of Difficulty for each fixture. The inverse of the slope of the Fitts' Law curve is computed by linear regression and listed in Table 2.9. This information is also represented in graphical form in Figure 2.8. Comparing each fixture tested, the two fixtures that resulted in the greatest overall increase in information processing capacity were fixtures 7 and 8, the only two to introduce auditory information. This analysis further supports the idea that using virtual fixtures to introduce abstract sensory information from multiple modalities is a viable method of enhancing operator performance.

2.6 CONCLUSIONS

The results of this study confirm that overlaying abstract perceptual information in the form of *virtual fixtures* on top of the sensory feedback from a remote environment can greatly enhance performance in telemanipulation tasks. Virtual fixtures composed of simple combinations of rigid impedance surfaces and abstract auditory information increased operator performance by up to 70% in a standard peg insertion task. It should be noted that a peg insertion task was chosen for this study because it offered an effective means of

quantifying operator performance and not because it was thought to be a particularly good application for virtual fixtures. Although virtual fixtures proved to be an effective means of improving peg insertion performance, it is likely that the fixtures would have had a much greater impact upon the performance of a more challenging telemanipulation task.

Because effective virtual fixtures were developed from very basic elements like rigid impedance surfaces and simple gradient fields, the development of an interactive *perceptual workstation* that allows an operator to build virtual fixtures from basic building blocks seems like a feasible endeavor. Such an interactive workstation could allow a teleoperator to develop powerful virtual fixtures upon first encounter with an unfamiliar task in an unstructured environment. Such a workstation might also allow an operator to quickly modify a fixture as task conditions change.

Critical comparison among the basic perceptual elements tested in this study along with the results of post testing interviews suggests that virtual fixtures enhance operator performance in the following ways:

1. Virtual fixtures simplify the operator's interaction with the remote workspace by restricting unnecessary kinematic freedoms and preventing unwanted actions from occurring. This simplification limits the alternatives the operator has to consider when executing the task and therefor reduces the information that needs to be processed. This simplification also allows the operator to relax criteria for success and failure by eliminating some modes of failure and by providing assistance in achieving some aspects of a successful task completion. What is more, virtual fixtures represent such task simplifications in perceptual terms which has been shown to be more effective than simple software constraints upon remote manipulator activity.

2. Post testing interviews suggest that virtual fixtures alter the operator's conceptualization of the remote task by embellishing the user's perceptual reality with abstract perceptual information. Without the aid of any virtual fixturing, the remote peg insertion task was primarily a visual operation. Thus, the operator was likely to conceptualize the peg insertion task as a manual procedure that "looks a certain way." With virtual fixtures, additional haptic and auditory information is provided to the operator. Such embellishments of the user's perceptual reality allows the operator to conceptualize the peg insertion task as not just "looking a certain way" but also as "feeling a certain way" and "sounding a certain way."

3. Virtual fixtures were shown to provide a localizing reference to the remote worksite by introducing haptic and auditory cues that can be coupled to the operator's kinesthetic sense of workspace geometry. This was demonstrated dramatically by test Fixture 5 which allowed a 20% performance improvement by providing only haptic localizing information to the operator's unused hand. It was also

hypothesised that all virtual fixtures tested acted to enhance an operator's sense of presence in the remote site by providing additional perceptual pathways through which afference and efferece can be related.

4. Virtual fixtures may displace the burden from taxed sensory modalities by distributing information through alternative sensory pathways. Without the aid of virtual fixtures, the peg insertion task primarily taxed the visual modality. With the fixtures, haptic and auditory sensations could also be relied upon. Had this task required additional use of the visual modality for scanning the scene or monitoring some other aspect of the task, the beneficial effect of displacing the burden from the visual modality would likely have been even more pronounced.

Chapter Three

3. Virtual Fixtures and Time-Delay Telepresence

This chapter continues the exploration of virtual fixturing by examining the effectiveness of perceptual overlays to enhance human performance when transmission delays are introduced into the telepresence link. Time delays as small as 200 ms result in perceptually significant misregistrations between an operator's kinesthetic sense-of-self and the sensory feedback reported from the remote environment. Such misregistrations between afferent and efferent perceptual cues distort the operator's perceptual reality and diminish the illusion of presence within the remote site. As a result, time delays in the telepresence interface cause significant degradations in operator performance. It has been observed that operators of time delayed telemanipulation systems are forced to adopt an less efficient "move-and-wait" strategy to compensate for such perceptual misregistrations during task execution [Ferrell, 1965; Ferrell 1966].

This chapter addresses whether or not perceptual overlays in the form of virtual fixtures can be used to reduce the performance degradation resulting from the introduction of time delays into a telepresence interface. Because perceptual overlays reside in the operator-space (i.e. they are completely independent of all information reflected from the remote environment), they are not effected by the time delays that distort the perceptual representation of the remote environment. Thus perceptual overlays act as real time perceptual cues in an otherwise time distorted perceptual reality. The remainder of this chapter describes an empirical study designed to test the hypothesis that enhancing an operator's perceptual reality through the introduction of real-time perceptual references can reduce the performance degradation resulting from time delays in the telepresence link.

3.1 INTRODUCTION

3.1.1 Telepresence and the Effect of Time Delay

The fundamental purpose of a telepresence system is to extend an operator's sensory-motor facilities and problem solving abilities to a remote environment [Pepper and Hightower, 1984]. Telepresence is achieved by projecting an operator's manipulatory dexterity to a remote site while reflecting sensory information from that environment so realistically that the operator feels present at the remote location [Akin et al., 1983]. The illusion of *presence* is dependent upon the fidelity and scope of the interface. As defined by Loomis [1993], a fundamental requirement of operator presence in a telepresence system is an understandable relation between afference and efference as was described in detail in Chapter 1. Such a perceptual requirement implies that a predictable correlation between the operator's movements (as perceived kinesthetically), and the corresponding activities in the remote site (as perceived visually, aurally, and haptically) is required for an effective telepresence interface [Held et al., 1991]. The introduction of time delay between master and slave distorts the correlation between kinesthetisia and sensory feedback from the remote environment, disrupts the illusion of presence, and ultimately degrades performance within the remote environment [Ellis, 1991]. Time delays as short as 100 ms have been found to degrade performance in telemanipulation tasks. Because destructive time delays are inherent to all telepresence systems in which the remote sites requires a round trip communication transmission time of more than 100 ms, performance degradation resulting from time delays is considered to be a fundemental research problem in telepresence laboratories. Many researchers have attempted to provide solutions for the performance degradation resulting from time delayed telemanipulation.

Ferrel [1965] was the first to demonstrate experimentally the performance degradation resulting from time delays in teleoperation. He found that operators confronted with time delay had a tendency to adopt an less efficient "*move and wait*" strategy. Thompson [1977] showed how task completion times were slowed by both time delay and reduction in degrees of constraint in part mating tasks. Since that time, much research has been devoted toward combating the effect of temporal delays through supervisory control and predictive displays. *Supervisory control* solves the time delay problem by simply removing the human from the control loop. Supervisory systems allow an operator to specify high-level commands to an autonomous manipulator in a remote environment [Ferrell and Sheridan, 1989]. Although this technique is valuable for tasks that can be performed by an autonomous manipulator, it does not solve the time delay problem for those cases when human presence is required. Predictive displays, on the other hand, combat the time delay problem while still maintaining some degree of operator presence in the remote environment. Rather than providing the operator with time delayed feedback from a remote worksite, a predictive display provides the operator with simulated sensory feedback generated through

simulated interaction with a computer model of the remote environment. The interaction with the computer model can occur in real time, while the real telemanipulation task actually occurs some time delay later. In theory, interaction with a highly accurate and perceptually rich virtual *model* of a remote environment would provide an operator with equivalent sensory feedback to actual telepresence in the *real* remote environment. In practice, computational requirements limit most predictive displays to sparse feedback from the simulated remote site consisting of a very simple visual representation [Bejczy, 1990; Hashimoto and Sheridan, 1986; Noyes and Sheridan, 1984]. In a few cases, limited use of other sensory modalities such as basic haptic cues have also been incorporated into predictive displays [Hannaford, 1989; Kotoku, 1992; Paul, Lindsay, and Sayers, 1992].

Although predictive displays have proven effective in combating performance degradation due to time delay, the use of such systems is greatly limited by the requirement for accurate modeling. For a predictive display to be effective, the remote manipulator and remote workspace must be modeled accurately enough that predictive interactions between the simulated slave manipulator and the simulated remote environment can allow an operator to effectively perform a task in the *real* worksite. As a result, the use of predictive displays is limited to simple tasks that are performed in well structured, well known environments. Even in situations in which the remote environment and telepresence task is suitably simple to allow for the implementation of an effective predictive display, a time consuming modeling effort is required up front which greatly limits the applicability of predictive displays.

An alternative or supplement to the implementation of supervisory control and predictive display systems is the use of *virtual fixtures* to reduce performance degradation due to time delay in remote manipulation tasks. Because the use of virtual fixtures does not involve modeling of the remote workspace, detailed previous knowledge of the remote site is not required. Thus, virtual fixturing is particularly promising for enhancing time delayed teleoperation tasks in unstructured, changing environments where other methods fail. Before describing an empirical study which demonstrates the effectiveness of virtual fixtures to reduce the performance degradation due to time delays, a brief clarification of virtual fixture concept is given.

3.1.2 Virtual Fixtures and Time Delay

As demonstrated in Chapter 2, the use of simple virtual fixtures greatly enhances operator performance in a telemanipulation task without time delay. The results revealed that operator performance was increased up to 70% when simple virtual fixtures were used as compared to the no-fixture case. The results suggested that virtual fixtures enhance performance by simplifying the conceptualization of the task, providing localizing references to the remote worksite, providing additional modalities through which

afference and efference could be related, and by reducing the demands on taxed sensory modalities by providing perceptual information through multiple sensory pathways. When time delays are introduced into teleoperation systems, information lag between the user space and the remote workspace degrades performance by distorting the correlation between the user's kinesthetic sense-of-self and the sensory feedback from the remote worksite. Such time distortions hinder user localization to the remote workspace and reduce the sensation of remote presence [Ellis, 1991]. Because of the perceptual problems that time delays introduce, it is believed that the beneficial effects previously shown to be provided by virtual fixtures in teleoperation without time delay will be even more pronounced when time delays are included.

Although virtual fixtures are overlaid on top of the reflection of the remote environment, they are completely *independent* of all sensory feedback from that remote site. Because virtual fixtures do not depend on information being fed back from the remote site, user interaction with fixturing is immune from the temporal distortions that result from communication delays. Virtual fixtures therefor provide the operator with real-time reference points that are overlaid on top of the time distorted representation of the remote site. It is hypothesized that the introduction of such real-time sensory cues can counter-act the effect that time delay has in weakening the correlation between the operator's kinesthetic sense-of-self and the sensory feedback of remote manipulator dynamics. Virtual fixtures are expected to enhance user localization to the remote worksite and strengthen the overall illusion of remote presence by providing sensory information that can be interacted with in real time.

Consider a simple line drawing task, as used as an example in Chapter 1 but now include time delay in the scenario. Imagine a human operator is controlling a dexterous remote manipulator that is holding a pencil in its gripper. The operator is given the simple task of drawing a straight line between two points on a piece of paper in the remote environment. Because drawing a straight line requires constant visual supervision and hand/eye coordination, even a small delay between the operator's imparted hand motion and visual feedback from the remote site can greatly impair operator performance. Such tasks often result in operator instability when the user tries to compensate for deviations from the desired trajectory with real-time manual adjustments while being provided with time delayed sensory feedback. Lines drawn under such conditions often reveal a characteristic oscillation resulting from the operator's futile attempts to react to past events.

Predictive displays have been implemented to combat these sorts of time delay problems by providing the operator with sensory feedback from a real-time computer simulation of the remote workspace. For the above line drawing telemanipulation task, the user of a predictive display would perform the task by interacting with a graphic computer model the remote pencil and paper workspace rather than with the actual time delayed feedback from the remote site. If the model of the remote manipulator

holding the pencil is accurate, the user's interaction with the simulation will cause the real pencil to draw a straight line in the real remote environment at some time delay after the user finishes manipulating the predictive model. However, if a small error existed in the simulation, the execution in the remote site could deviate greatly from the desired trajectory without the operator being aware. Although the use of a predictive display is a viable solution to time delay problems in many simple telemanipulation tasks, the requirement for accurate modeling of the remote manipulator and remote workspace restricts the use of predictive displays to structured environments where the workspace and task are well known in advance.

A proposed alternative or supplement to the use of predictive displays, which does not require a detailed modeling of the remote workspace, is the implementation of virtual fixturing. Rather than simulating a real-time model of the remote pencil and paper worksite, the virtual fixture approach simply provides the user with a real-time model of a virtual *ruler* that can be overlaid on top of the reflection of the real remote environment. Such a ruler would be a generic virtual fixture, chosen from a toolbox of fixtures, that could be overlaid on top an arbitrary remote environment. The use of a virtual ruler in such a line drawing task would likely shift the user's sensory dependency from the time delayed visual feedback reflected from the remote site to the real-time haptic feedback provided by the fixture. As a result, the performance degradation due to time delay would likely be reduced. What is more, the kinematic constraints and haptic cues provided by the ruler-like virtual fixture would likely result in the user drawing a much straighter line than if no fixture had been used. Of course this simple example of virtual fixture implementation can be extended to much more complex multidimensional tasks as well as more abstract virtual fixturing. It is in these more complex situations where the application of virtual fixturing is most appropriate.

3.2 Experimental Testing

3.2.1 Fitts' Law Performance Test

To quantify teleoperator performance in a remote manual task, a Fitts' Law paradigm was chosen because of its general acceptance as a robust measure of human performance. The same peg insertion task was used in this set of tests as was used in the initial testing described in Chapter 2. For a complete description of the Fitts' Law Paradigm and Testing Procedure see Section 2.1.

3.2.2 Experimental Hardware

To implement a standardized Fitts' Law performance test in a telepresence environment, the following hardware subsystems were used: Peg Insertion Task Board, MBA Exoskeletal Master / Merlin

Robot Arm Slave, Monocular Vision System, and Virtual Fixture Board. As shown in Figure 3.1, the system is divided into two physically separate parts: the *remote environment* and the *operator space*. For more details on this set up, refer to Chapter 2, Section 2.2.

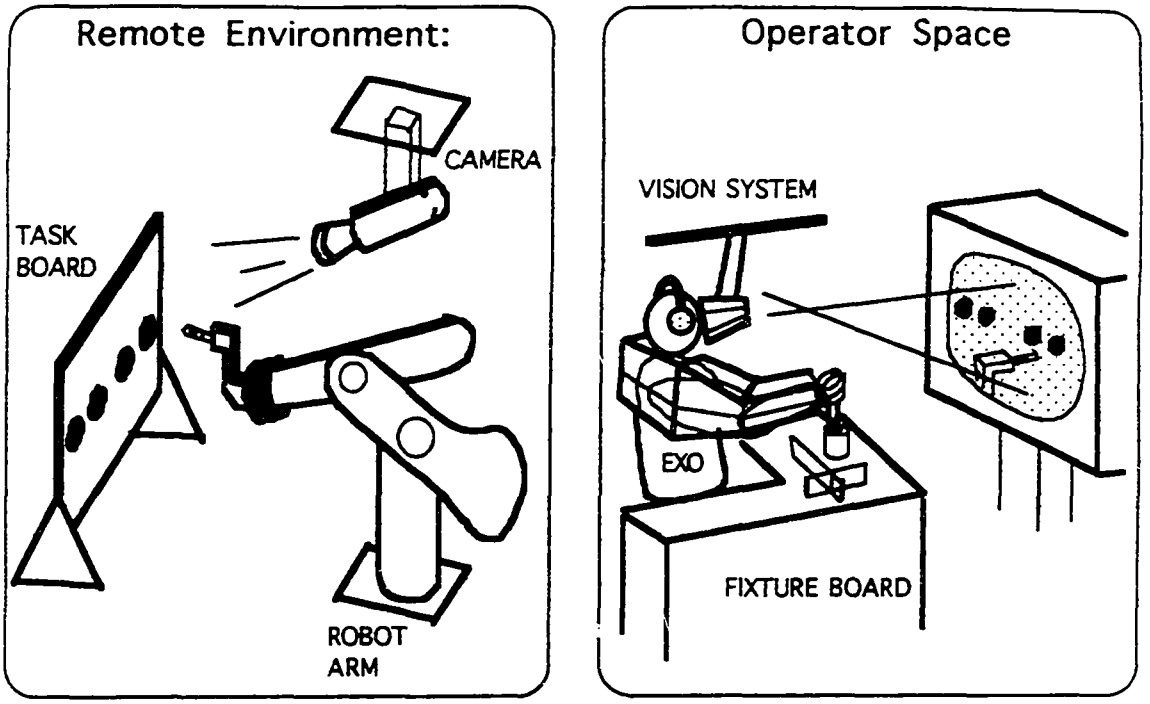

Figure 3.1: Complete Telepresence System developed to implement the testing of teleoperator performance in a standardized peg insertion task with and without the aid of virtual fixtures.

The remote robot arm is slaved to the right hand position of the exoskeleton. Thus, when the subject moves his right hand to interact with the image of the task board, the end-effector on the slave arm follows. Because the subject cannot see his own hand when looking into the vision system but does see the remote robot end-effector in the position where he feels his hand to be, a sense of presence within the remote environment is created. The subjects also wear a set of stereo headphones for use in implementing auditory perceptual overlays. Auditory surfaces and fields can be modeled on the control computer and interacted with by the user. The complete system provides a powerful testbed for projecting subjects into a remote environment and overlaying haptic and auditory information on top of the reflected percepts.

3.3 EXPERIMENTAL DESIGN & PROCEDURE

3.3.1 Experimental Design

The goal of this study was to investigate the effect of overlaid virtual fixtures has upon performance in a standard teleoperated task with time delay. As described in Chapter 2, a previous experiment compared eight virtual fixture designs in a telemanipulation task without time delay. For this experiment, three of the eight previously tested fixtures were used for comparison to a no-fixture control case. Each of the three fixture configurations was tested with no delay, 250 ms delay, and 450 ms latency between master and slave. Figure 3.2 shows a schematic representation of each of the fixtures tested. (NOTE: The fixtures are labeled with the same numbers given in the previous chapter for reference purposes.) These fixtures are shown overlaid on top of the fixture board as they are perceived by the subjects. Although these fixtures are represented graphically in this figure, they are only perceived haptically by the subjects. While technical limitations prevented the visual representations of virtual fixtures in this study, such an experimental simplification should not dimish the fact that presenting virtual fixtures with visual representation is a viable application of perceptual overlays and an important topic for future investigation.

The three virtual fixtures tested are composed of simple combinations of planar surfaces which are perceived haptically as rigid wall sensations by the users. As shown in Figure 3.2, virtual Fixture 4 is composed of three intersecting planar surfaces. One surface is simply a horizontal plane oriented like a table top in the workspace and aligned so that contact with the rigid surface will result in vertical alignment of the operator's hand with the holes in the peg board. The second plane runs diagonally across the first plane in such a way that interaction with this rigid surface will guide the subject's hand motion across the face of the board and inward towards the target hole. The third plane runs parallel with the target hole such that contact with the plane will halt hand motion when the peg passes in front of the target hole.

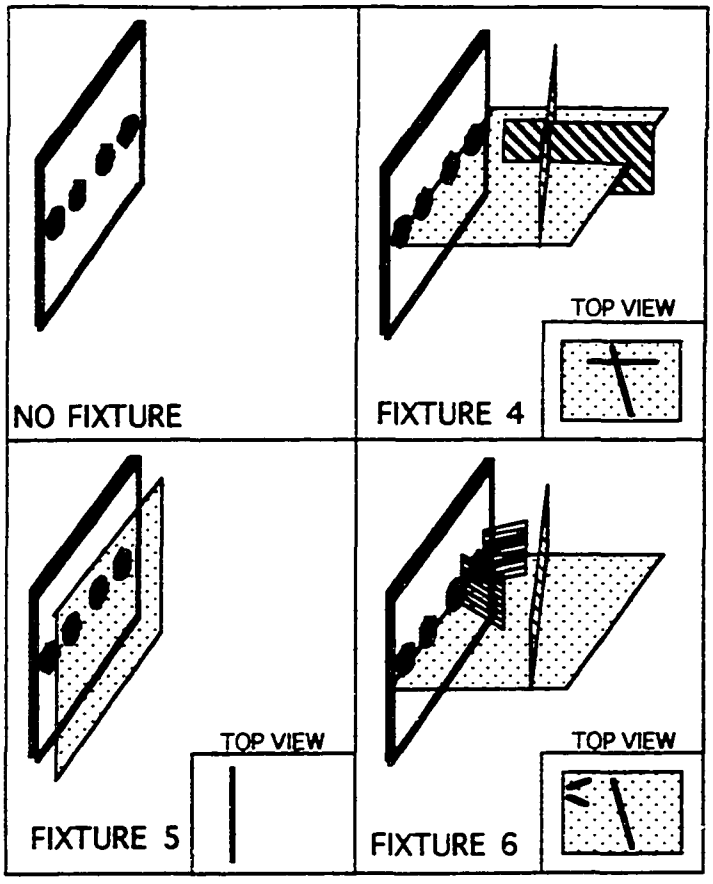

Figure 3.2: Virtual fixtures shown overlaid onto workspace as they are perceived by subjects. Also shown are top views, looking down from above the task board. These fixtures are shown graphically here, but they are perceived haptically by subjects.

Virtual Fixture 6 is similar to virtual Fixture 4 except that the surface which stops hand motion in front of the target hole is removed and two diagonal surfaces are added in its place. These two diagonal surfaces are positioned and oriented to assist in the fine positioning of the peg into the target hole.

Virtual Fixture 5 is very different from the two fixtures presented thus far in both its geometry and implementation. Fixture 5 is a rigid surface oriented parallel to the task board and located approximately 0.5" in front of the board. This fixture was *not* designed for interaction with the subject's right hand (the hand that performed the peg insertion task), but rather was designed for interaction with the subject's *unused* left hand. Subjects were told to place their left hand upon the surface while performing the task with the right hand. The intent of this fixture is to isolate the effect of localization on performance in time delayed teleoperation tasks. A previous study of virtual fixtures in teleoperation without time delay, as described in Chapter 1, suggested that interaction with virtual fixtures provides localizing information to the user which can enhance understanding of the remote workspace. It was hypothesized that the localization provided by virtual fixtures allows for better correlation of the operator's kinesthetic sense of hand position to the geometry of the remote site. It is desired now to investigate if enhanced localization alone can reduce the detrimental effects that time delay has upon telemanipulation. Because the other fixtures to be tested in this study act to guide the operator's hand motion as well as to provide localizing references to the remote workspace, it would be impossible to isolate the effects of the localizing information alone. Fixture 5 was therefore designed; it only interacts with the unused hand and thus can only have a localizing effect upon operator performance.

3.3.2 Experimental Protocol

A series of tests was run to evaluate subject performance using each fixture configuration for each of three time delay periods: no delay, 250 ms delay, and 450 ms delay. Operator performance was recorded during test periods which included 12 practice and 36 timed peg insertion trials. A single trial consisted of moving a peg from a designated start hole to a designated target hole. The holes were referenced by numerals (3, 4, 5 or 6) located above each hole. Two start hole/target hole pairs were studied in these tests: a 16 cm motion from hole 3 to hole 6 and a 4 cm motion from hole 5 to hole 6. The 36 trial period was divided into three groups of 12 trials. Each of these groups required the subject to perform the insertion task using a different peg size. After the completion of each group of 12 trials, subjects were instructed to rest for approximately two minutes while one of three possible peg sizes was installed in the robot end-effector.

Reproduced with permission of the copyright owner.

The use of two motion amplitudes (4 cm and 16 cm) and three peg sizes (0.75 cm, 0.98 cm, 1.50 cm diameter pegs) allowed for the testing of insertion trials with six different task difficulties. Table 3.1 shows all combinations of peg size and motion amplitude and lists the Index of Difficulty for each task as dictated by Fitts' Law [Fitts, 1954]. The order in which the three peg sizes were presented to each subject was randomized to ensure that mental and physical fatigue had similar effects on all task difficulties.

Table 3.1: *Index of Difficulties (ID)* shown for all combinations of peg size and peg motion amplitudes as predicted by Fitts' Law.

	Motion Amplitude	
Diameter	5 to 6 (4 cm)	3 to 6 (16cm)
0.75 cm	**2.68** bits/response	**4.68** bits/response
0.98 cm	**2.98** bits/response	**4.98** bits/response
1.50 cm	**4.00** bits/response	**6.00** bits/response

Subjects were instructed to begin each test period with a 3 to 6 peg insertion trial. Subjects would then perform a 5 to 6 peg insertion trial, then a 3 to 6 peg insertion trial, and cycle in that manner throughout the 36 trials. Subjects performed each insertion trial as described previously in Chapter 2, Section 2.3.2. Subjects were tested over 6 experimental sessions, each lasting 45 to 60 minutes. To minimize the effects of fatigue and boredom, no subject completed more than one experimental session during a single day of testing. The first two experimental sessions included three test periods of 36 trials. The first period of the session was performed with no fixture and no delay, the second was performed with no fixture and 250 ms delay, and the third was performed with no fixture and 450 ms delay. These initial 216 trials were treated only as practice during which the subjects familiarized themselves with the use of the exoskeleton, Merlin robot arm, vision system, and fixture table. It was also intended that subjects get experience with the effect that time delay has upon telemanipulation.

It was found that by the end of the second practice session, all subjects had sufficiently learned the task that variability in movement times for trials of the same difficulty had fallen below 20% for every subject, with a mean coefficient of variability of 15% for all subjects. Once learning had stabilized, subjects were sequentially tested using each of the test fixtures for each of three time delay scenarios. A new fixture configuration was tested during each of the next four daily sessions. Table 3.2 lists all of the

daily sessions along with the fixtures tested during that session. The order of the last four daily sessions was randomized for each subject to reduce the effect that learning might have upon performance results.

Table 3.2: *Testing Schedule:* All four tests shown along with corresponding daily sessions and fixtures tested on that day. Presentation order of tests I,II, III, and IV were randomized.

TEST	SESSION	FIXTURES STUDIED
Practice	day 1	No Fixture: no delay 250ms delay 450ms delay
Practice	day 2	No Fixture: no delay 250ms delay 450ms delay
Test I	day 3	No Fixture: no delay 250ms delay 450ms delay
Test II	day 4	Fixture 5: no delay 250ms delay 450ms delay
Test III	day 5	Fixture 4: no delay 250ms delay 450ms delay
Test IV	day 6	Fixture 6: no delay 250ms delay 450ms delay

3.4 RESULTS

3.4.1 Raw Data

The following section presents the raw results of peg insertion trials for each time delay period for comparison among each of the virtual fixture configurations tested. The basic quantity presented is the recorded movement times for peg insertion trials. The consistancy of this raw data is reflected in the Coefficient of Variation (C.O.V.) which has been computed across the trials of the same index of difficulty performed with the same time delay and same virtual fixture configuration. The C.O.V. is therefore a measure of subject consistency in performing a particular task with a particular latency and using a particular virtual fixture. Tables 3.3 through 3.6 list the mean movement time results across all six subjects for each index of difficulty and each time delay configuration. Table 3.7 lists the coefficients of variation in movement times for each time delay and fixture configuration tested.

TABLE 3.3: *Mean Movement Times* for each Index of Difficulty are shown for six subjects using NO FIXTURE to perform peg insertion tasks with no delay, 250 ms delay, and 450 ms delay. Also computed is the mean movement time across all trials (across task difficulty).

TEST I: no fixture trials

	INDEX OF DIFFICULTY						
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean
no delay	833ms	927ms	991ms	1617ms	1600ms	1666ms	1272ms
250ms	1144ms	1246ms	1374ms	2123ms	2142ms	2362ms	1731ms
450ms	1165ms	1284ms	1500ms	2140ms	2329ms	2636ms	1842ms

TABLE 3.4: *Mean Movement Times* for each Index of Difficulty are shown for six subjects using FIXTURE 5 to perform peg insertion tasks with no delay, 250 ms delay, and 450 ms delay. Also computed is the mean movement time across all trials (across task difficulty).

TEST II: fixture 5 trials

	INDEX OF DIFFICULTY						
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean
no delay	656ms	662ms	807ms	1448ms	1396ms	1497ms	1077ms
250ms	819ms	882ms	975ms	1683ms	1814ms	1892ms	1344ms
450ms	846ms	1003ms	1181ms	1941ms	2015ms	2273ms	1543ms

TABLE 3.5: *Mean Movement Times* for each Index of Difficulty are shown for six subjects using FIXTURE 4 to perform peg insertion tasks with no delay, 250 ms delay, and 450 ms delay. Also computed is the mean movement time across all trials (across task difficulty).

TEST III: fixture 4 trials

	INDEX OF DIFFICULTY						
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean
no delay	679ms	682ms	791ms	1375ms	1467ms	1506ms	1083ms
250ms	762ms	757ms	956ms	1473ms	1547ms	1625ms	1187ms
450ms	696ms	596ms	741ms	1518ms	1513ms	1540ms	1100ms

TABLE 3.6: *Mean Movement Times* for each Index of Difficulty are shown for six subjects using FIXTURE 6 to perform peg insertion tasks with no delay, 250 ms delay, and 450 ms delay. Also computed is the mean movement time across all trials (across task difficulty).

TEST IV: fixture 6 trials

	INDEX OF DIFFICULTY						
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean
no delay	641ms	691ms	811ms	1266ms	1361ms	1461ms	1039ms
250ms	618ms	669ms	717ms	1489ms	1417ms	1579ms	1082ms
450ms	635ms	591ms	721ms	1365ms	1499ms	1539ms	1058ms

TABLE 3.7: *Mean Coefficient of Variation* for all subjects performing tasks of the same difficulty with each fixture as shown. Also shown is the overall C.O.V. which indicates the overall consistency of the results.

	DELAY	C.O.V.
NO FIXTURE	No delay	14% ±3.3
	250 ms	14% ±3.2
	450 ms	18% ±3.9
FIXTURE 5	No delay	13% ±1.3
	250 ms	13% ±1.7
	450 ms	17% ±2.7
FIXTURE 4	No delay	8% ±1.5
	250 ms	9% ±1.0
	450 ms	14% ±1.8
FIXTURE 6	No delay	10% ±2.8
	250 ms	14% ±2.0
	450 ms	12% ±3.0
All Trials	MEAN	13% ±2.4

3.4.2 Processed Data

In order to quantify the drop in operator performance resulting from the introduction of time delays into the telemanipulation task in a way that could be easily compared among the various virtual fixture configurations tested, percentage increase in movement time was computed as a unitless measure of the effect that time delay had upon operator performance. Thus performance degradation was quantified as the percentage increase in movement time for peg insertion trials with either 250 ms and 450 ms delays as compared to the trials with no time delay. Tables 3.8 through 3.11 list these percentage increases. This information is summarized graphically in Figure 3.4.

TABLE 3.8: *Performance Degradation Due to Time Delays* for all six subjects performing task using NO FIXTURE. Performance degradation computed as percentage increase in movement times for peg insertion trials with 250 ms and 450 ms delays as compared to the trials with no time delay.

TEST I: no fixture trials

	INDEX OF DIFFICULTY							
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean	stdev
250ms	37.3 %	34.4 %	38.6 %	31.3 %	33.9 %	41.8 %	36 %	±10.8
450ms	39.9 %	38.5 %	51.4 %	32.3 %	45.6 %	58.2 %	44 %	±17.4

TABLE 3.9: *Performance Degradation Due to Time Delays* for all six subjects performing task using FIXTURE 5. Performance degradation computed as percentage increase in movement times for peg insertion trials with 250 ms and 450 ms delays as compared to the trials with no time delay.

TEST II: fixture 5 trials

	INDEX OF DIFFICULTY							
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean	stdev
250ms	24.9 %	33.3 %	20.8 %	16.2 %	30.0 %	26.4 %	25 %	± 9.5
450ms	28.9 %	51.5 %	46.3 %	34.0 %	44.4 %	51.8 %	43 %	±17.8

TABLE 3.10: *Performance Degradation Due to Time Delays* for all six subjects performing task using FIXTURE 4. Performance degradation computed as percentage increase in movement times for peg insertion trials with 250 ms and 450 ms delays as compared to the trials with no time delay.

TEST III: fixture 4 trials

	INDEX OF DIFFICULTY							
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean	stdev
250ms	12.2 %	11.1 %	20.8 %	7.1 %	5.5 %	7.9 %	10 %	± 2.9
450ms	2.5 %	-12.6 %	-6.3 %	10.4 %	3.2 %	2.3 %	2 %	± 6.8

TABLE 3.11: *Performance Degradation Due to Time Delays* for all six subjects performing task using FIXTURE 6. Performance degradation computed as percentage increase in movement times for peg insertion trials with 250 ms and 450 ms delays as compared to the trials with no time delay.

TEST IV: fixture 6 trials

	INDEX OF DIFFICULTY							
DELAY	2.68	2.98	4.00	4.68	4.98	6.00	mean	stdev
250 ms	-3.7 %	-3.3 %	-11.6 %	17.6 %	4.13 %	8.1 %	4 %	± 10.1
450 ms	-1.0 %	-14.6 %	-11.1 %	7.9 %	10.1 %	5.4 %	2 %	± 8.9

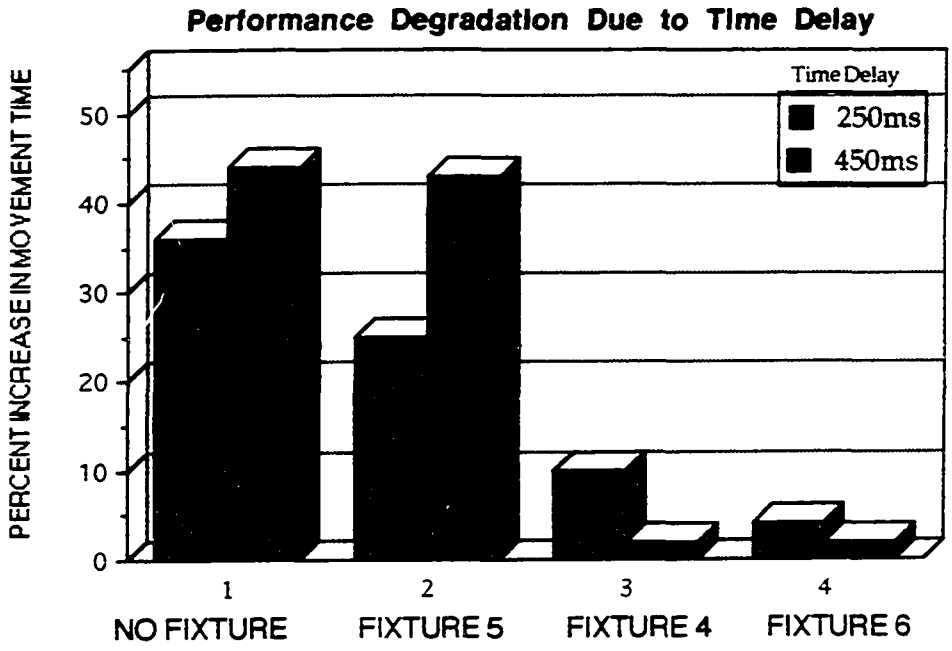

Figure 3.3: Performance Degradation Due to Time Delay for Each of the four Virtual Fixture Configurations Tested. Performance Degradation Computed as Mean Percentage Increase in Movement Time for Peg Insertion Trials.

3.5 DISCUSSION

The following section addresses the implications of performance degradation due to time delay and investigates the effect that virtual fixturing has upon operator performance in time delayed teleoperation. The objective of this analysis is, first, to quantify the effect that time delay has upon operator performance in a controlled telemanipulation task and, second, to determine if virtual fixtures can be used to effectively reduce the resulting drops in operator performance. Performance degradation has been computed as the percentage increase in movement time for a given task when a particular time delay is introduced to the interface (as compared to the no-delay case).

Before comparing operator performance, the issue of repeatability and reliability of the results is addressed. As shown in Table 3.7, the coefficient of variation in movement times for trials with a given time delay and fixture configuration is under 18% for each combination of test parameters. The mean coefficient of variation is 13% across all test parameters. Such consistency suggests that the results of the experimental sessions are highly reliable and implies that subjects were sufficiently trained during the practice session.

3.5.1 Peg Insertion Trials Without Virtual Fixture

As listed in Table 3.8 and shown in Figure 3.3, the time required to successfully complete a peg insertion trial was found to increase dramatically when time delays were introduced into the telemanipulation interface. For trials (without fixtures) with small time delays of 250 ms and 450 ms, movement times increased, on average, by 36% and 44% respectively. This dramatic increase in movement times resulting from the introduction of small time delays replicates earlier work that showed that operator performance in even the most simple telemanipulation tasks is highly susceptible to degradation due to time delays. Assuming an ideal telemanipulation system with no time lag due to information processing and computation, time delays of 250 ms and 450 ms would correspond to communication transmission lags that would occur if the operator and the remote site were separated by distances as small as 37000 km and 67500 km respectively. Such distances would be typical when controlling a remote manipulator in geosynchronous earth orbit from a command station on the ground [Johnsen and Corliss, 1967]. The results of this study confirm that in order to achieve efficient telemanipulation in earth orbit, some means of reducing performance degradation due to communication transmission lags is required. In a non-ideal telemanipulation system, latencies due to information processing and computation would compound the total time delay and further reduce performance.

3.5.2 Peg Insertion Trials With Virtual Fixtures

As listed in Table 3.9 through Table 3.11 and shown in Figure 3.3, the increases in the movement time for successfully completion of peg insertion trials when virtual fixtures were used was significantly reduced as compared to unaided trials. Turning first to peg insertion trials using virtual Fixture 5, Figure 3.2 shows that this fixture is composed of a single rigid impedance plane which is parallel to the task board surface and located approximately 0.5" in front of the board. Whereas the other two fixtures tested were designed for interaction with the operator's right hand (the hand that manipulates the peg), Fixture 5 is designed for interaction only with the operator's unused left hand. Subjects place the palm of their left hand upon the surface while they perform the task with their right hand. This fixture was tested to isolate the effect of localization upon performance. As shown in Figure 3.3, the increases in movement times associated with time delays of 250ms and 450ms were 25% and 43% respectively. Comparing this to the 36% and 44% increases in unaided trials, we see that the use of Fixture 5 resulted in a small drop in performance degradation for the 250ms delay trials but no drop in performance degradation for the 450 ms delay trials. This result suggests that this fixture did little to reduce the performance degradation due to time delay. This is probably because Fixture 5 did not interact with the operator's active hand (it interacted with the operator's passive hand which was not manipulating the peg) and therefor could not provide additional sensory pathways through which afference and efference could be related. As a result, the presentation of Fixture 5 did not provide a significant enhancement to the operator's time distorted perceptual reality.

Turning next to peg insertion trials performed with the aid of virtual Fixture 4, we recall that Fixture 4 is composed of three intersecting rigid surfaces as shown in Figure 3.2. Unlike Fixture 5, Fixture 4 interacts with the user's active hand (i.e. the hand which is manipulating the peg). Fixture 4 is composed of a combination of the following three haptic percepts: 1) a horizontal surface oriented like a table top in the workspace and aligned so that contact with the surface will result in vertical alignment of the operator's hand with the holes in the peg board; 2) a planar surface diagonally across the first such that interaction with this surface will guide a subject's hand across the face of the board and inward towards the target hole; 3) a surface parallel with the target hole such that contact with the plane will halt hand motion when the peg is in front of the target hole. The fixture surfaces are perceived as overlaid haptic sensations only; subjects cannot see the fixtures when performing the task. It should also be noted that interaction with the overlaid haptic sensations provides real-time sensory feedback to the operator even though all interaction with the remote site is time distorted.

When provided with the real-time haptic information from virtual Fixture 4, performance degradation in the time delayed telemanipulation task was significantly less than in the no-fixture trials. As listed in Table 3.10 and shown in Figure 3.3, peg insertion trials using virtual Fixture 4 resulted in movement time increases of only 10% and 2% for introduced time delays of 250ms and 450ms respectively. Comparing this result to the 36% and 44% movement time increases for the no-fixture trials, it is clear that the use of virtual Fixture 4 has greatly reduced performance degradation resulting from time delays in the telemanipulation interface. In fact, for the longer delay of 450ms, the use of virtual Fixture 4 results in movement times that reflect no drop in operator performance as compared to the no-delay trials. The fact that less performance degradation was recorded for the 450ms trials than was recorded for the 250ms trials implies that subjects had an easier time taking advantage of the real-time sensory information provided by the virtual fixture in the trials with the longer delay. In post testing interviews, subjects reported finding it easier to focus their concentration on the real-time haptic sensations provided by the virtual fixtures (ie. focus less attention on the time distorted feedback from the remote site) when the delay period was longer. In psychophysical terms, when the latency is short, the time delayed sensory information competes with real-time sensory information provided by the perceptual overlay and thus *masks* the effects. The shorter the delay, the more the real-time and delayed sensations compete, and the greater the masking. Thus when implementing virtual fixtures, there is some perceptual advantage to having a longer time delay in the telepresence link. This surprising result suggests a counter intuitive principle of design for perception (which would not have been derived through simple physical modeling or control system analysis) that increasing the time delay in a telepresence link might actually be used by designers to enhance operator performance.

Turning next to peg insertion trials performed with the aid of virtual Fixture 6, we recall that virtual Fixture 6 is similar to virtual Fixture 4 except in that the surface which stopped hand motion in front of the target hole is removed and two diagonal surfaces are added in its place. These two diagonal surfaces are positioned and oriented so as to assist in the fine positioning of the peg into the target hole. When subjects made use of the real-time haptic embellishment provided by virtual Fixture 6, the performance degradation in the time delayed telemanipulation was significantly less than in the no-fixture trials. As listed in Table 3.11 and shown in Figure 3.3, peg insertion trials using virtual Fixture 6 resulted in movement time increases of only 4% and 2% for time delays of 250ms and 450ms respectively. Since neither of these results reflects significant increases in movement times, it is clear that when using virtual Fixture 6, subjects performed just as well with 250ms and 450ms time delays as they did in the no-time-delay trials. This result demonstrates that virtual fixtures can be used effectively to reduce, and in this case actually eliminate, the performance degradation due to time delay.

In post testing interviews subjects were asked why they felt their performance was so much better with Fixtures 4 and 6 than with no fixture in time delay trials. Subjects reported that without the aid of virtual fixturing they had to adopt a move-and-wait strategy to allow the feedback from the remote site to catch up with their kinesthetic sense of peg position. They reported moving in front of the target hole and then waiting until the visual feedback from the remote site confirmed that their position was correct. With the aid of virtual fixturing, subjects were able to use the overlaid real-time haptic cues in place of the delayed visual cues and perform the task without having to wait for visual confirmation of every action. Subjects reported having a greater sence of presence in the remote site because they could relate their kinesthetic sense of hand position to the force feedback provided by the fixtures even through the visual feedback was time distorted. This result lends support to the theory proposed by Loomis [1993] which states that the illusion of presence is dependent upon the extent to which afference and efference can be related.

3.5.3 Characteristic Trajectory Paths For Peg Insertion Trials

Each subject was asked to perform a small number of peg insertion trials during which time their hand trajectories were recorded to disk. Such trajectories were captured for each subject performing peg insertion trials including every combination of virtual fixture configuration and time delay period. This data was collected and plotted in order to gain qualitative insights into the effect of virtual fixtures on time delayed telemanipulation. Although many trajectories were recorded, a few representative plots were chosen for display in this section. Figures 3.4, 3.5 and 3.6 show characteristic hand trajectories for a single subject performing a peg insertion task at each of the three time delay configurations (0ms, 250ms, and 450ms plots). For each of these three plots, two curves are drawn: a position trajectory and a corresponding velocity profile. The top half of each plot shows the operator's hand trajectory for a single peg insertion trial. The plot depicts horizontal hand motion in which the peg starts at hole 3, is pulled away from the task board, is moved across the board, and is finally inserted into target hole 6. The horizontal axis corresponds to motion across the face of the task board and the vertical axis corresponds to distance away from the task board. The bottom half of each plot shows the velocity profile that corresponds hand trajectory ploted directly above it.

Figure 3.54shows sample trajectories for a subject performing a peg insertion task from hole 3 to hole 6 with no time delay, 250ms delay, and 450ms delay. For the no-delay trial, the position and velocity plots reveal that the subject pulled the peg out of the board, moved it quickly across the board, and then slightly hesitated before inserting the peg. This hesitation is revealed by a slight dip in the velocity profile right before peg insertion occurs. For the 250ms delay trial shown in Figure 3.4, we see a similar pattern,

but the slight hesitation before peg insertion has turned into a substantial pause. This pause is reflected as a dip in the velocity profile in which the velocity actually drops to almost zero. Such a pause confirms the move-and-wait strategy that subjects reported during post testing interviews. Stalls in hand motion are used to let the delayed visual feedback catch up to the operator's kinesthetic impression of hand position, thus allowing afference to match efference for a brief instant. Turning next to the 450ms delay trial shown in Figure 3.4, the single pause is replaced by a a series of three drops in hand velocity. Although these three pauses are difficult to see on the velocity graph because of overlapping lines, the raw data does reveal three distinct instances in which the operator's hand velocity dips to almost zero during target acquisition. This heavy reliance upon a move-and-wait insertion strategy resulted in a pronounced performance degradation as compared to the no-delay trial. If we compare the overall velocity plots of the three time delay scenarios, it is evident that velocity profiles get progressively flatter (slower) as the time delay increases. Such drops in overall velocity confirms graphically the performance degradation due to the addition of time delay into the telemanipulation task.

Figure 3.4

Sample trajectory plots for NO FIXTURE trials shown for each time delay period

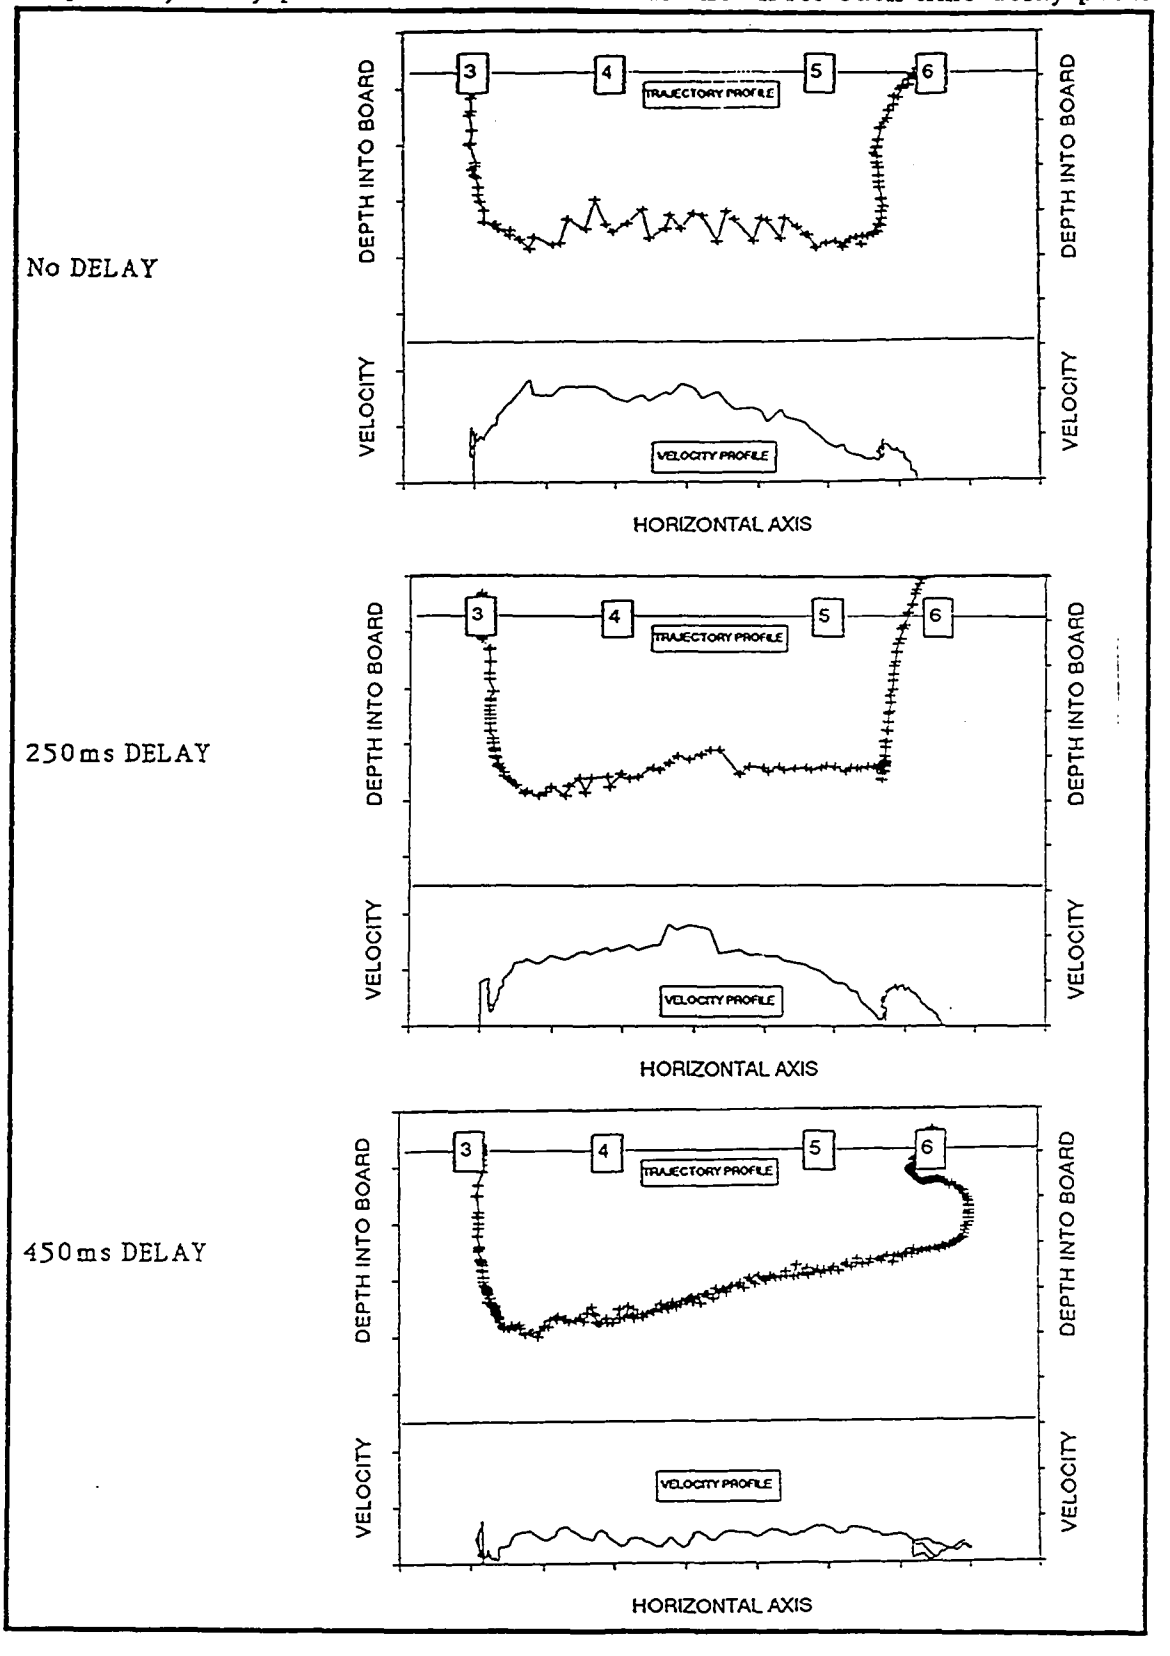

Figure 3.5 shows trajectory plots for a single subject using virtual Fixture 4 to perform the peg insertion trials with each time delay period. A top-view line drawing of this fixture's haptic surfaces is drawn on each graph to indicate where haptic interactions occurred during task execution. Comparing the three plots in Figure 3.5 to the three no-fixture plots in Figure 3.4, we observe that all three trajectories using Fixture 4 depict smoother paths with higher velocity profiles. In addition, we find that the the dips in velocity before peg insertion are less pronounced in all three time delay trials which implemented virtual Fixture 4. This result suggests that subjects were able to use haptic sensations as confirmation of kinetic operations and did not have to rely as heavily on visual confirmation through move-and-wait strategies. The implications of these haptic cues are clearly revealed by comparing the hand trajectories for the 450ms delay trials. When no fixture is used, a characteristic oscillation occurs during the target aquisition stage of the insertion task. This oscillation is the result of the operators attempts at reacting to past events. The trajectories recorded for Fixture 4 show no such oscillations, suggesting that subjects can make use of the the real time haptic cues in place of time delayed visual information. What is more, the characteristic flattening (slowing) of velocity profile with increased time delay observed in Figure 3.4 for no-fixture trials is much less pronounced in Figure 3.5 for trials using Fixture 4. This graphically represents the reduction in time delay induced performance degradation when the virtual Fixture 4 was used.

Figure 3.6 shows trajectory plots for a subject using Fixture 6 to perform the peg insertion trials with each time delay period. A line drawing of this fixture's haptic surfaces is drawn on each graph to indicate where haptic interactions occurred during task execution. Comparing the three plots in Figure 3.6 to the three no-fixture plots in Figure 3.4, we observe that the dips in velocity profile which represented pauses in operator trajectory are not present in the Fixture 6 trials. This suggests that, using virtual Fixture 6, subjects were able to use the haptic sensations in place of real time visual confirmation and execute the task without using any move-and-wait techniques. Comparing the shape of the velocity graphs for each time delay period, we find no flattening in the velocity profiles with increasing time delay. In fact, when using virtual Fixture 6, the velocity profiles actually bulge with increasing time delay, reflecting progressively faster performance rather than the performance drop suggested by no-fixture trajectory profiles. These trajectory plots clearly demonstrates that, for a given subject, the use of a virtual fixture can eliminate the pronounced detrimental effects that time delay imparts on telemanipulation.

Figure 3.5

Sample trajectory plots for FIXTURE 4 trials shown for each time delay period

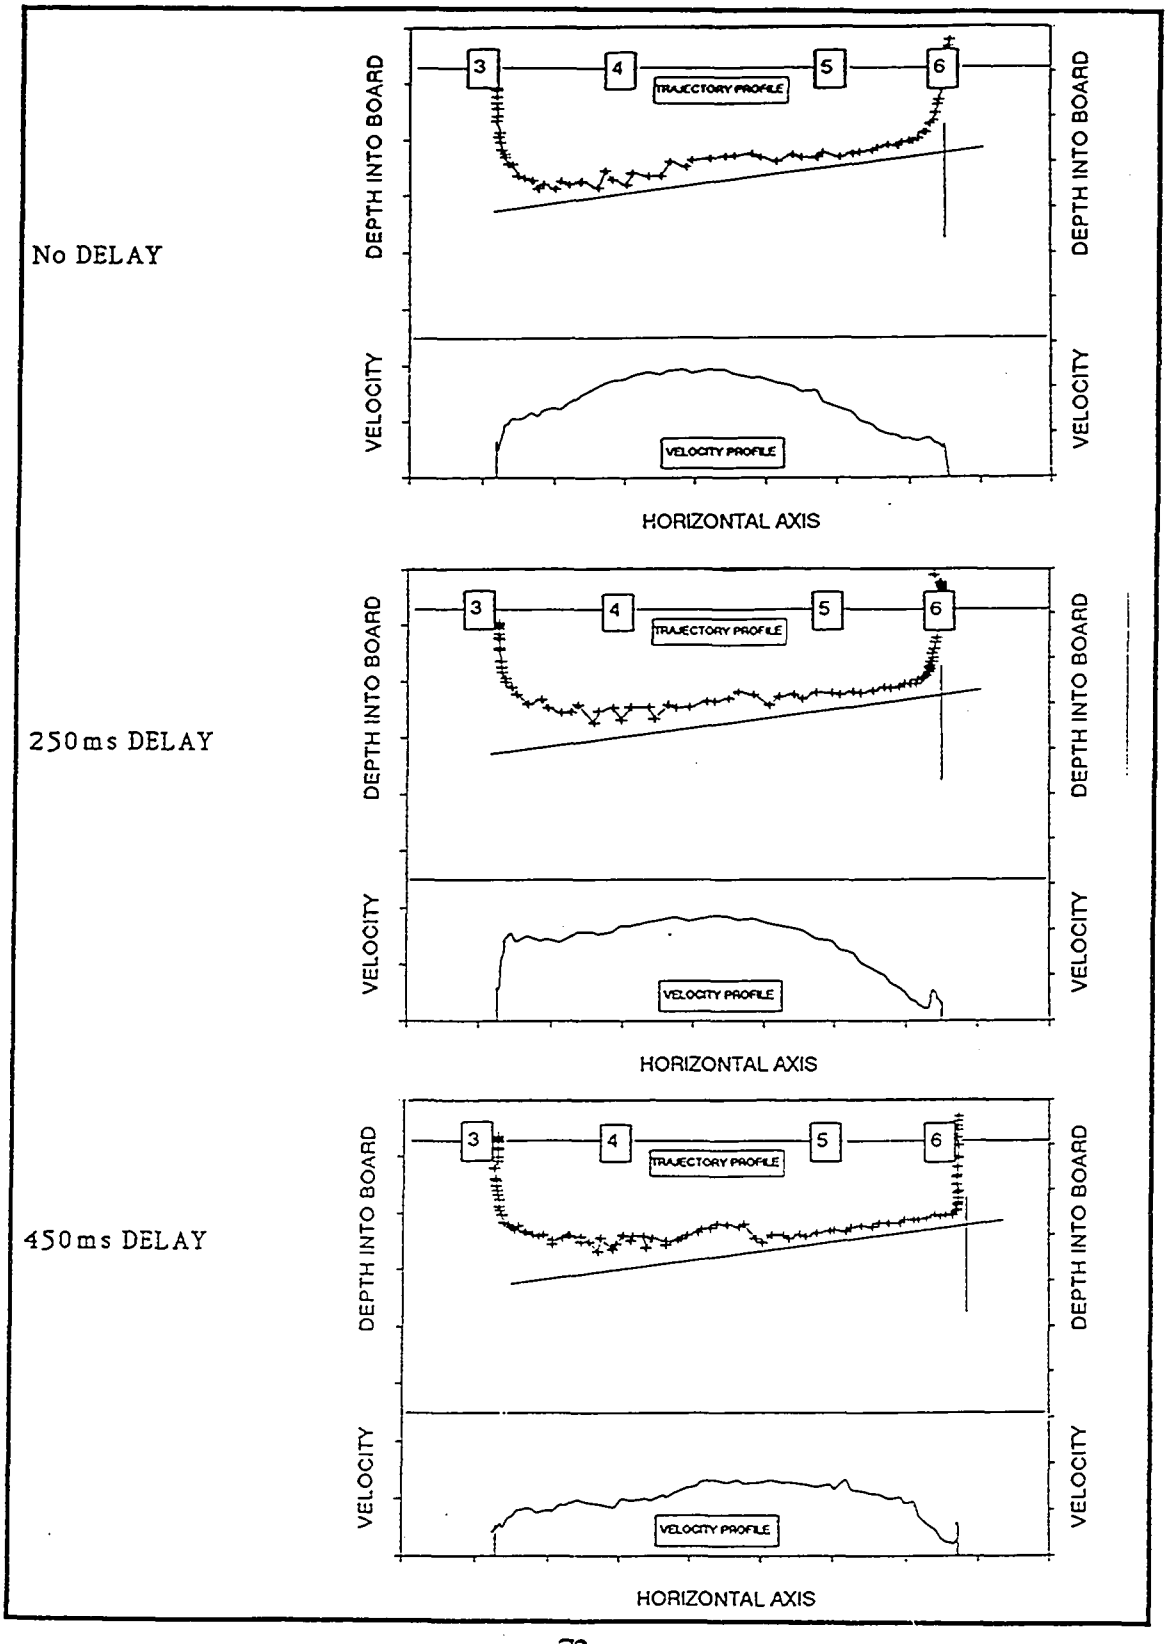

Figure 3.6
Sample trajectory plots for FIXTURE 6 trials shown for each time delay period

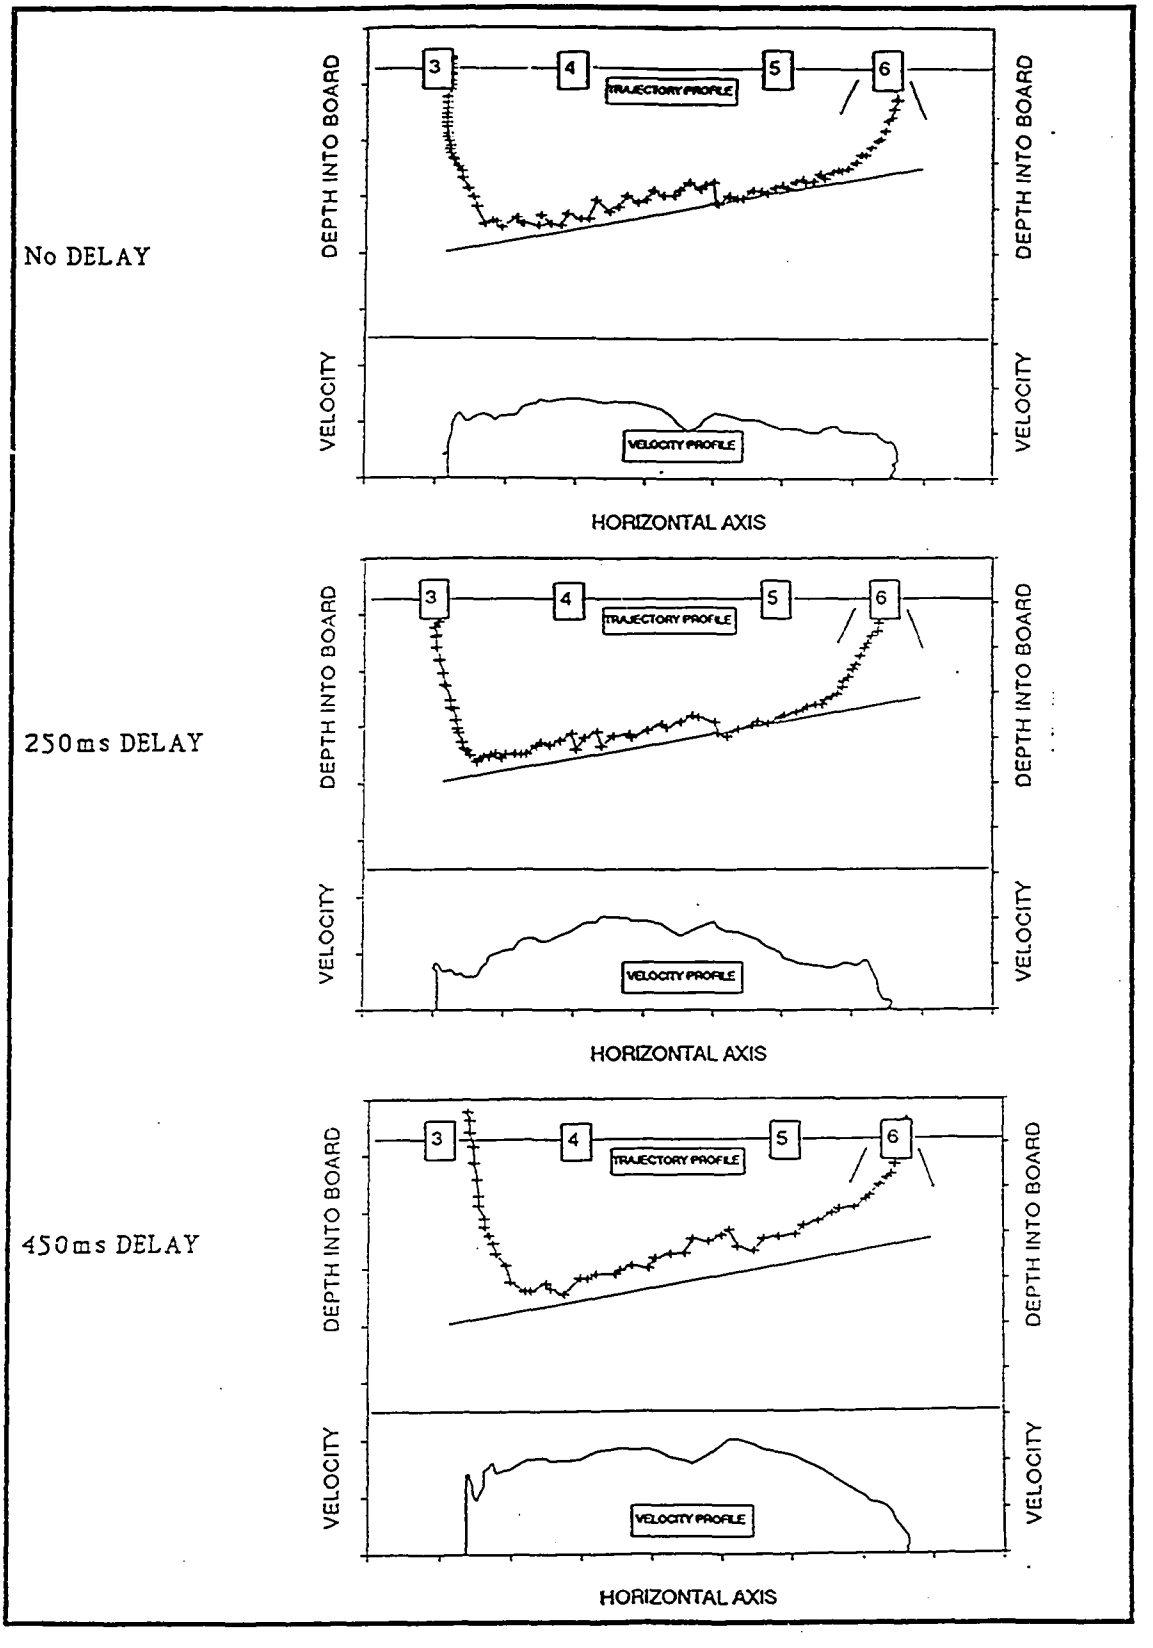

3.6 CONCLUSIONS

The results of this study confirm that enhancing an operator's perceptual reality with abstract sensory information in the form of *virtual fixtures* can greatly dimishish the performance degradation that is induced by time delays in a telepresence link. Without such overlays, the operators perceptual reality is temportally distorted and therefor fails to reflect a lawful relationship between afference and efference. With the use of virtual fixtures, real-time perceptual cues are introduced which allow operator's to compensate for the temportal distortions reflected from the remote site. Without the use of any virtual fixtures, 36% and 44% drops in operator performance (increases in movement times) were recorded for time delays of 250ms and 450ms respectively. Such large drops in operator performance relfect the profound effect that small temporal distortions in perceptual representation can have upon human performance. During post testing interviews, subjects reported being forced to adopt a move-and-wait strategy to deal with the introduced time delays. Such move-and-wait task execution was confirmed by recorded trajectory plots.

The use of virtual fixtures composed of simple combinations of haptic surfaces were found to greatly reduce, and in some cases completely eliminate, the performance degradations resulting from time delays in the telepresence link. Subjects reported that they were able to rely upon overlaid haptic precepts as real-time perceptual cues in place of time delayed visual cues. As a result, subjects could abandon inefficient move-and-wait techniques and perform the task more naturally. The absence of move-and-wait motion was confirmed in trajectory plots of peg insertion trials with virtual fixtures.

Because effective virtual fixtures were developed from basic haptic elements such as rigid planar surfaces, the development of an interactive *perceptual workstation* which allows operators to build their own virtual fixtures from basic building blocks seems like a feasible endeavor for time delayed telepresence tasks. Such an interactive workstation could allow a teleoperator to develop powerful virtual fixtures upon first encounter with an unfamiliar telemanipulation task to be performed in a time delay situation. Although haptic cues used in isolation have proven effective in this study, it is likely that a perceptual workstation which supported additional modalities including visual and aural perceptual overlays, would have an even greater impact upon time delayed telemanipulation.

Chapter Four

"perception does not operate with the mechanical faithfulness of a camera, which records everything impartially... Seeing means grasping some outstanding features of objects; the blueness of the sky, the curve of a swan's neck, the rectangularity of the book, the sheen of a piece of metal, the straightness of a cigarette..."

- Rudolf Arnheim
Art and Visual Perception

---

4. DESIGN FOR PERCEPTION & THE RIGID WALL PERCEPT

4.1.0 Overview

Thus far we have explored the notion of virtual fixturing using real physical surfaces to represent ideal haptic perceptual overlays. The next step is to extend such studies to the more general case in which the sensory information is actually computer generated. Before testing the viability of computer generated virtual fixtures to enhance operator performance in a controlled experimental paradigm (as will be discussed in Chapter 5), it is worth taking a brief detour to explore the design issues which underlie the generation of such virtual information. This chapter investigates the display of virtual haptic information and asks the fundamental question: How does one represent a *realistic* haptic percept through a force reflecting interface. Just by asking this question we reveal an array of slippery issues relating to realism, reality, and human perception.

In Chapter 1, the notion of *Design for Perception* was introduced and applied to the design of virtual environments. A clear distinction was made between *physical reality* and *perceptual reality* and the phenomenon of *distal attribution* was described which relates these realms in our consciousness. *Design for perception* was presented in Chapter 1 as a methodology in which virtual percepts are generated by representing the perceptual rather than physical content. It was proposed that if an acceptable perceptual representation of a virtual environment was produced and displayed to a user, the process of distal attribution would infer physical realism even if the representation was not physically accurate. Thus the design requirements for the generation of an effective virtual environments need not include the *accurate*

representation physical phenomenon, but must include the presentation of *appropriate* perceptual information.

This chapter revisits the notion of *design for perception*, but rather than addressing the global issues related to the design of effective virtual environments, the following sections focus on lower level issues related to the design of a single percept. The particular percept chosen for this *case study* is the haptic representation of a rigid wall. This percept was chosen as a highly effective testbed for *design for perception* because force reflecting systems are notorious for producing very poor representations of rigid surfaces. This case study is particularly insightful because all documented attempts at generating haptic models of rigid surfaces have been approached through pure physical modeling. Thus this study is the first attempt a modeling the *sensation* produced by a rigid wall rather than just modeling the physics of a rigid wall.

The following pages introduce the basic principles of haptic display and describe a common haptic entity to be referred to as a "rigid-surface-contact". The methodology of *design for perception* is then applied to the generation of such a percept. Finally, an empirical study is described which assesses the effectiveness of this design strategy. Although this chapter is in many ways a detour from the central notion of virtual fixtures, such an exploration into the design of haptic percepts strengthens the perceptual foundation on which the concept of virtual fixturing is built.

4.1.1 The State of the Art in Haptic Display

While the state of the art in visual and auditory information display for virtual environments has been advancing by leaps and bounds in recent years, the representation of virtual force information is still in its early infancy. Hindered by hardware limitations and safety considerations, most force reflection systems are unable to convincingly represent even the most basic of haptic percepts. What is probably the most frequently distorted sensory percept attempted by virtual force reflection systems is that of a rigid surface contact. Although most of the tactual interactions we encounter in our daily lives involve contact with rigid surfaces, most force reflecting systems are unable to realistically reproduce such a percept. Virtual rigid surface contacts are often described as "mushy", "sticky", or "bouncy" by even the most forgiving users. Jex [1991], in reporting on informal "rules of thumb" derived from experience with high-performance force reflecting aircraft simulators, suggested that the ability to produce a convincing rigid wall is a primary requirement of any general purpose haptic interface. Because of the basic importance of the realistic display of haptic rigid surfaces, this study critically examines the "rigid-surface-contact" percept and attempts to develop design guidelines for the convincing generation of such haptic sensations.

4.1.2 The Technology of Force Reflection

A force reflecting virtual environment system consists of force reflecting interface hardware and a computation engine. The interface hardware typically consists of a mechanical linkage in the form of a joystick or exoskeleton which couples the human operator to a source of mechanical power—either electromagnetic, electrohydraulic, or electropneumatic actuators. The computation engine governs the dynamic behavior of the interface hardware as a function of measurements from interface transducers and algorithms that describe the virtual models to be simulated. A distinguishing aspect of the display of haptic virtual information as compared to the display of visual or aural virtual information is that the same body part and interface hardware is used for both input and output between the human and the virtual environment. Consequently, haptic information transfer is affected not only by the processing capacity of the computation engine and the comprehensiveness of the virtual model, but by the controlled dynamics of the interface linkage and the human limb itself. Because the physical properties of the coupled interface is highly dependent upon the changing physical properties human limb, designers who chose to produce haptic simulations through physical modeling face a particularly challenging analytical task.

Nevertheless, the implementation of virtual haptic surfaces has been demonstrated in a variety of experimental studies [Kilpatrick, 1976; Winey, 1981; Adelstein, 1989; Minsky et al, 1990; Fasse, 1993]. Acknowledged potential shortcomings of rigid wall simulations include: high frequency vibration [Ouh-young, 1990]; low frequency instability [Grafing, 1992]; excessive compliance [Kilpatrick, 1976; Winey, 1981; Ouh-young et al., 1988]; and stickiness [Adelstein, 1991]. Stability is one problem of haptic wall implementation which has been addressed analytically [Ouh-young, 1990; Chin, 1991; Colgate et al., 1993].

4.1.3 Design for Perception applied to the Design of Haptic Percepts

Although previous research has addressed the generation of virtual rigid surface contacts, such work has approached the problem by attempting to generate accurate *physical models* of a rigid surface within the hardware constraints of force reflecting devices. Although such investigations of stability and dynamics response are useful in defining limitations of force reflection hardware systems, such physical representations offer little insight into the perceptual aspects of presenting virtual force information. Recent studies have modeled virtual surfaces as linear spring-damper systems because such systems are convenient to analyze from a dynamics and controls perspective [Adelstein, 1989; Minsky et al, 1990; Colgate et al., 1993]. These studies strive to develop virtual models which dynamically *behave* as close to a real rigid surface contact as hardware constraints will allow. Although it is reasonable to assume that a

system which reproduces the *exact* dynamic behavior of a rigid surface contact will be *perceived* by a user exactly like a real rigid surface, it is not necessarily true that a physical model which *almost* behaves like a rigid surface will be perceived as something *almost rigid*. Since hardware limitations prevent even the state-of-the-art force reflecting systems from accurately modeling stiffness values on the order of magnitude of a real rigid surface, attempting to generate *similar* dynamic behaviors may not be the best approach to developing the most believable virtual percepts.

As defined by the *design for perception* methodology introduced in Chapter 1, it is the belief of this author that the construction of virtual force information from a *perceptual* rather than physical perspective may result in the generation of more believable percepts within the hardware limitations of force reflecting systems. For example, it may be the case that a virtual model of a stiff linear spring is mathematically similar to a real rigid-surface-contact, but is perceptually very different when the stiffness value is significantly smaller than that of a real rigid surface. On the other hand, a linear viscous damper may be physically very different than a real rigid-surface-contact, but it might contain strong perceptual similarities. Although dynamic analysis has led many researchers to models of rigid surfaces composed of a linear spring that is as stiff as their hardware will allow plus a small amount of added damping to enhance the stability characteristics, there is no reason to believe that such models contain particularly convincing perceptual content.

The question remains, how does one model a sensory percept such as a rigid-surface-contact in terms of its perceptual content rather than physical parameters? To answer this question we turn to the *perceptual psychology of art* and gain insights from a field which devotes itself to effective representations of the perceptual world. Rudolph Arnheim, in his book *Art and Visual Perception* [?], writes that perception starts with the grasping of a salient features and even suggests that "features are the primary data of perception." He goes on to say that "a few outstanding features not only determine the identity of an object, but also make it appear as a complete, integrated pattern." Thus, like an artist composing a painting, the designer of a haptic percept may compose the desired effect by combining salient perceptual features in a balanced composition. Unfortunately, the important perceptual features which represent haptic percepts are not as well known as those which represent visual information. Thus, for a designer of virtual environments to synthesize a haptic percept from perceptual foundations, some process of *perceptual analysis* must be developed to reveal the salient perceptual features.

4.1.4 Perceptual Psychology and the Design of a Virtual Percept

When doing a perceptual analysis of virtual sensory percepts, it is convenient to define terms such as *proximal stimulus*, *distal stimulus*, and the *perceptual hypothesis* to refer to the different stages of sensory perception. A proximal stimulus is often defined as the sensory information falling upon a receptor. A proximal stimulus may be an image falling upon the retina, a force imposed upon a muscle spindle, or a sound wave disturbing the basilar membrane of the cochlea. A distal stimulus, on the other hand, is the distant source of such sensory information. It is the reflective surface from which the visual image emanates, the weight which when lifted stretches the muscle spindles, or the bird from which the sound wave was born. Although we interact with an environment of distal stimuli, we as human beings only have access to proximal stimuli. Thus the act of perception is often described as the transduction of a proximal stimulus coupled with the *guessing* of what distal stimulus most likely caused the sensation. This act of inference is often called the perceptual hypothesis and results in the generation of an internal representation of the outside world known as a percept. As long as the perceptual system is presented with enough salient sensory information in a proximal stimulus, a correct perceptual hypothesis can be made and an appropriate internal model of the actual distal stimulus will be created [Levine, 1981].

For example, a proximal stimulus might be the image of a cube falling upon your retina as shown schematically in figure 4.1. Your perceptual system extracts the salient information such as edges, angles, shading, depth cues, and size cues from the proximal stimulus and then infers that the distal stimulus is a small cube located across the room at a particular location and orientation. It is important to remember that we do not have direct access to actual cube in physical world, but rather infer its existence through a perceptual hypothesis drawn from highly abstracted perceptual information. Once such a perceptual hypothesis is made and the distal stimulus is inferred, an appropriate mental model of the physical world is constructed. It is this mental model of the physical world with which we interact and not the actual physical world.

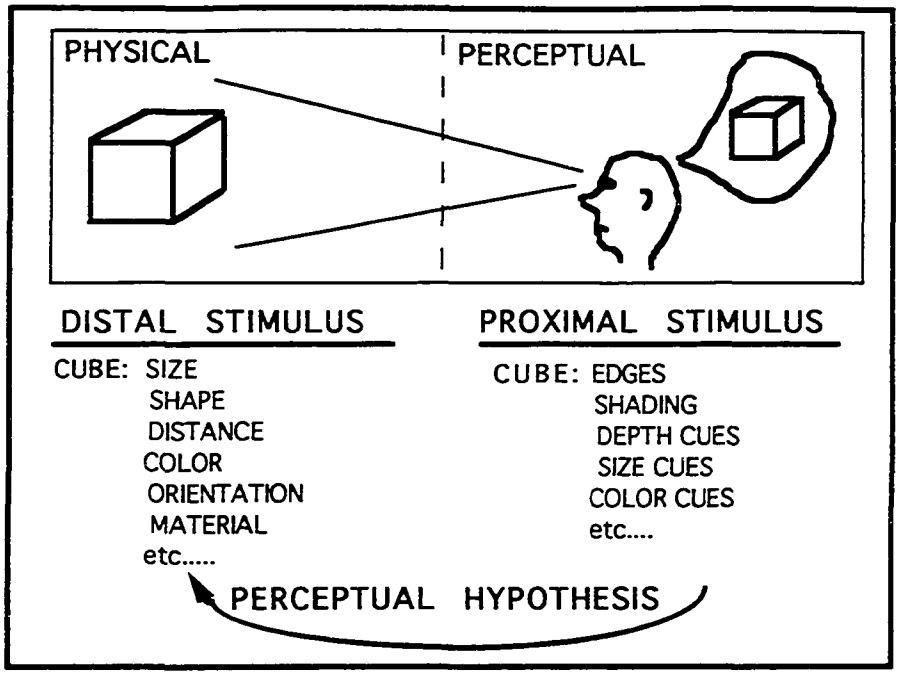

Figure 4.1: Perceptual Hypothesis: infer Physical from Perceptual

Although the above example implies that a simple relation exists between a given proximal stimulus and the corresponding distal stimulus which it describes, the act of perception is far from such a clean and simple process. Consider the proximal stimulus gathered from the cube in the above example. It is likely that the sensory information reaching your retina is very different from the sensory information gathered during previous instance when you viewed that same cube. The lighting conditions may have changed, your viewing location may have changed, it may even be a cube you have never actually seen before. Nevertheless you identify the object as a cube and build an internal percept. Our ability to draw the appropriate perceptual hypothesis despite changes in viewing conditions is called *perceptual constancy* and is important in allowing us to generate a robust internal model of the outside world. The fact that perceptual constancy can occur implies that some sensory information is critical for the identification and generation of a percept, while other information may not be as important. Because sensory perception is a complex process of inference based on certain features and not others, the key to designing an effective and efficient virtual percept is to ascertain which perceptual features are vital to the generation of percept in question.

Let us first consider the representation and presentation of visual information since it is the most familiar sensory modality for which people design. Whether our visual system is presented with a

photograph of a cube or a rough sketch of a cube, the image is likely to be identified as a cube and an appropriate internal perceptual model will be generated. In a sense, the photographic representation is analogous to physical modeling of the distal stimulus while the sketch representation is analogous to perceptual modeling of the proximal stimulus. Although a sketch contains much less sensory information than a photograph, the *sketch artist* is skilled at providing only the appropriate sensory features that assure desired perceptual analysis of the image. A good sketch can often be a more effective representation of sensory information than is a poor photograph. We can extend the analogy of the *sketch* and the *photograph* to virtual haptic sensations produced by force reflecting systems. Rather than producing a physically accurate (i.e., "photographic") representation of a haptic sensation, a *perceptual designer* could "sketch" haptic sensations by combining only those appropriate perceptual features which make up a desired percept. Such an approach may be more effective than "photographic" dynamic modeling of a haptic sensation, particularly in cases where force reflecting equipment lacks the fidelity to generate a realistic "photo".

Bringing together the principles of perceptual psychology introduced here and in Chapter 1, the process of perception can be represented schematically as shown below in Figure 4.2.

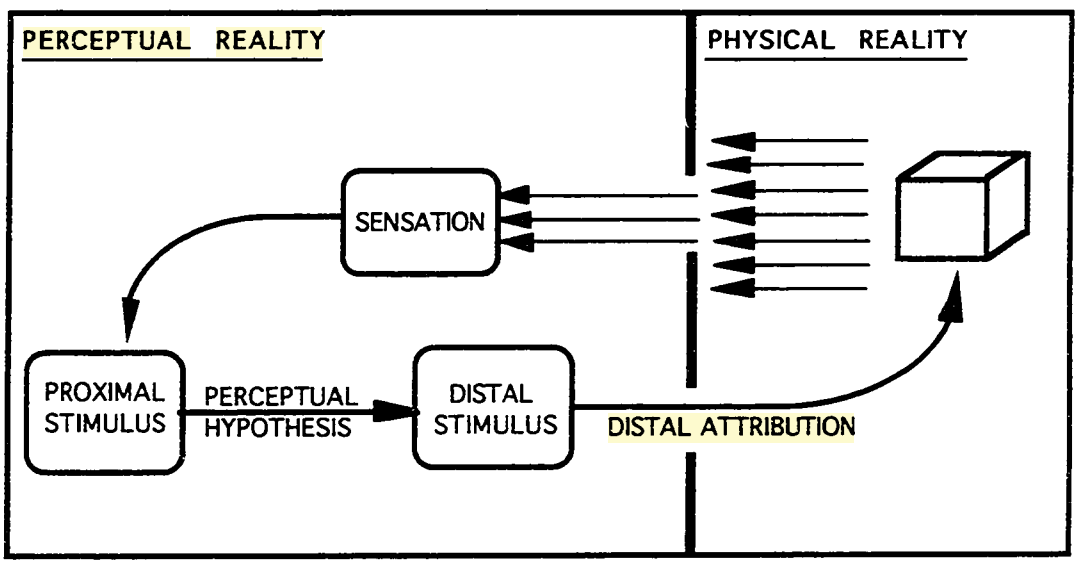

Figure 4.2: Schematic of the Process of Perception

The diagram demonstrates that *perceptual reality* and *physical reality* are separate and distinct realms and that access between them is restricted to limited pathways. As shown in the figure, the cube represents an arbitrary object in the physical surroundings. The arrows emanating from the cube represents the large amount of information which describes the physical properties of this object. Only a small band of this

information can be transduced by our sensory system. This is represented on the diagram by the three arrows which can actually enter our perceptual reality. Once such information enters our nervous system through our sensory receptors, it is further abstracted by the many levels of pre-processing which greatly reduces the information content before reaching the cognitive areas of our cerebral cortex. This collection of highly abstracted perceptual cues is what we call the *proximal stimulus*. The process of perception then proceeds by making a best guess as to what distal stimulus in the physical world must have caused the given proximal stimulus. This act of inference is known as the *perceptual hypothesis*. Once a distal stimulus has been inferred and a mental model of the object has been generated, the process of *distal attribution* occurs which relates the observer's mental model of the surroundings with the actual physical surroundings. In the end, the observer is given the illusion of directly experiencing a physical reality which he or she does not really have direct access to. In everyday life, this illusion allows people to directly experience the physical surroundings despite the many levels of perceptual abstraction which actually exists. In virtual reality and telepresence systems, this illusion allows users to experience synthetic environments as if they were the direct physical surroundings despite the fact that many levels of computational and perceptual abstraction actually exists.

Having defined this simple schematic representation of the process of perception, the first question to be addressed by designers of virtual percepts is this: Where in this process should we *present* virtual information? The answer to this question is highly restricted by technology, for until direct neural methods of information presentation are technologically possible, the location of virtual information display must occur prior to the process of sensation. With that restriction, the next important question to be addressed by designers of virtual percepts is this: Where in the process of perception should we *model* virtual information? The traditional approach is to model physical reality (i.e., produce the best possible representation of the many arrows that emanate from the cube in Figure 4.2). The problem with this approach is that hardware constraints often do not allow all such physical representations to be accurately represented. Without knowing which pieces of information are perceptually important (i.e., which arrows make it through to the perceptual side of the diagram) it is difficult for a designer to efficiently model physical reality and ensure that a viable percept will be generated. In addition to neglecting the perceptually important features, a designer may waste computational processing on the representation of physical features which are not perceptually important (arrows which do not make it through). Thus it seems that the most efficient and effective design strategy for the representation of virtual percepts is to synthesize the perceptual rather than physical content. Such a design process is shown schematically in Figure 4.3. As shown in the figure, the virtual model is generated in terms of a perceptual representation of the proximal stimulus. The cognitive processes of perception can then be relied upon to relate this perceptual representation with physical realism.

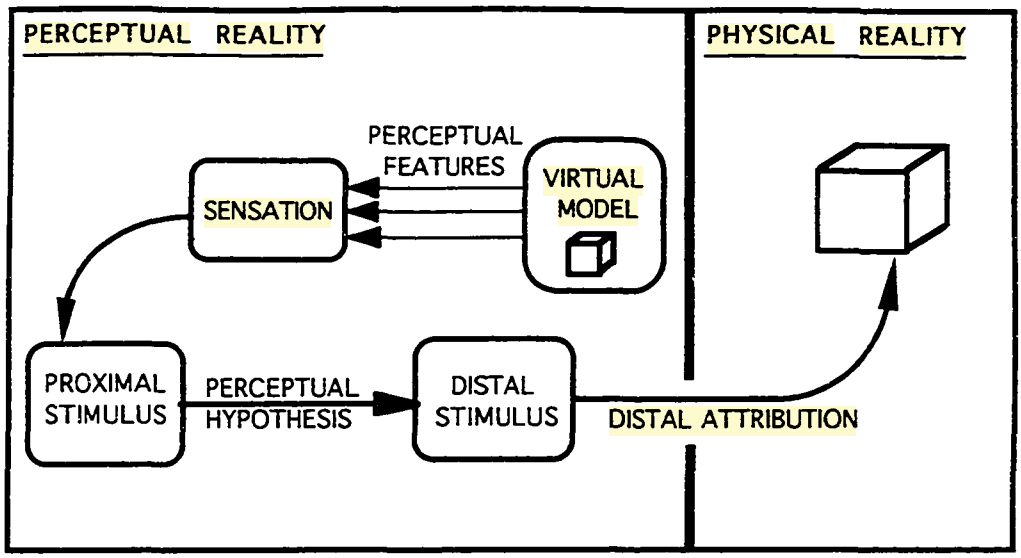

Figure 4.3: Generation, Presentation and Perception of a Virtual Percept

4.1.5 The Perceptual Analysis of a Rigid Surface Percept

When developing a virtual representation of a sensory percept, such as the sensation produced when encountering a rigid wall, the methodology of *design for perception* suggests little need for producing an accurate physical model of the *distal stimulus*. Rather, *design for perception* prescribes generating a perceptually adequate model of *proximal stimulus*. When such a model of the proximal stimulus is provided, the user should have little trouble making the correct perceptual hypothesis and inferring the appropriate distal stimulus (despite the fact that it was not directly represented by the virtual model). Although such a procedure is likely to produce an efficient representation of a virtual percept, it does pose some unique design challenges. The primary challenge is that in order to produce a *perceptually adequate* representation of a virtual percept, a strong understanding of the perceptual qualities of the percept being modeled is required. Such a perceptual understanding might be acquired through experience and training as is true of the visual artist learning to sketch the salient perceptual features of a visual scene. Unfortunately, when it comes to designing haptic percepts, cultural background and methods of formal training do not exist to guide and educate the perceptual designer to the effective modes of representation.

The first step in the perceptual design of a virtual rigid-surface-contact sensation was the development of an effective *decomposition* of the percept into its salient perceptual features. Because no methodology currently exists for performing such a perceptual analysis of a sensory percept, an initial attempt was made to perform a perceptual decomposition through observation and introspection. The

author spent many hours interacting with a two degree of freedom (dof) force reflecting joystick, gaining insight into the feel of a wide variety of virtual haptic representations. The author compared how the haptic sensations associated with simple virtual elements, such as springs and dampers, compared with the feel of a real rigid surface. The goal of such an exploration was to isolate distinct and independent perceptual qualities of the real rigid wall and reveal how to reproduce such perceptual qualities through simple virtual models.

4.1.6 Experimental Hardware for Haptic Display

Virtual models of simple physical elements such as springs and dampers were implemented using one degree of freedom (dof) of a high performance two axis force reflecting joystick built by Adelstein at MIT and shown in Figure 4.4 [Adelstein, 1989; Adelstein and Rosen, 1992]. Each axis of the joystick is powered by disk armature permanent magnet motors and is instrumented with optical encoders to sense position, tachometers for velocity, as well as accelerometers and an interface force transducer. As configured, the motors can produce continuous forces up to a maximum of 20N with zero cogging and negligible force ripple from DC up to 58 Hz (the first structural mode for the axis used in these experiments) at the joystick handle. The force threshold (i.e. starting friction) of the uncompensated joystick is 1 N. The joystick handle's passive inertia corresponds to a mass of 0.35 kg at the hand. In these experiments, the joystick is operated under purely digital control through an A/D and D/A card with DMA on an ISA bus Intel 486DX-50 based personal computer. The digital update rate for the control algorithms used in these experiments exceeded 10 KHz.

Figure 4.4: Two Degree of Freedom Force Reflecting Joystick

Using this hardware, virtual haptic sensations were implemented so that a user could grab the handle of the joystick and move it to the left until the handle encountered the virtual wall model. Such an arrangement allowed for both dynamic and static interaction with the simulated surfaces. For example, a virtual model of a pure linear spring was implemented so that when a user moved the handle far enough to the left, the simulated spring was encountered and the force reflecting joystick applied an opposing force proportional to compression of the spring. Such an implementation of a virtual linear spring element is shown schematically below in Figure 4.5. The figure depicts the joystick handle as it is treated computationally with respect to the virtual spring. The resistance force felt by the user is produced by the motors in the joystick base.

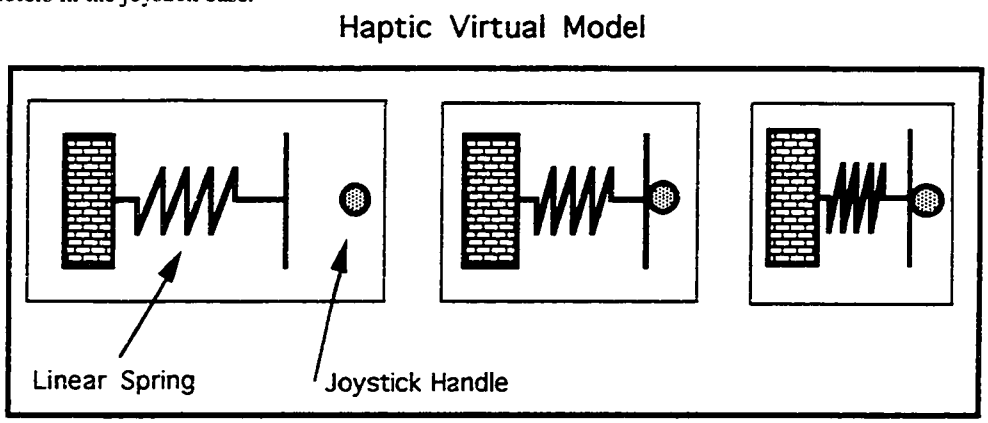

Figure 4.5: Schematic of Virtual Model of Haptic Surface Contact

Starting with simple virtual models such as linear springs and linear viscous dampers and expanding to nonlinear and more abstract elements, basic perceptual qualities of the virtual sensations were exhaustively compared to that of a real rigid surface contact. The real rigid surface contact was implemented by a physical hard stop that could be placed in the path of the joystick in the same location where the virtual models were presented. After extensive comparisons of many virtual models, the following decomposition of the rigid wall contact was hypothesized.

4.1.7 Perceptual Decomposition of a Rigid Wall Percept

The rigid wall percept was found to effectively decompose into at least three perceptual parts: the initial dynamic contact with the surface, the quasi-static interaction with the hard surface, and the final dynamic release from the surface. Each of these stages was observed to have distinct and independent perceptual qualities which can be described simply as the initial contact *crispness*, the rigid surface *hardness*,

and final release *cleanness*. Informal interviews with of a number of casual test subjects revealed that when such properties were varied, they were described with a range of verbal labels as listed in Table 4.1.

Table 4.1: Perceptual Decomposition of Rigid Wall Percept

Mode	Perceptual Content	Range of Descriptions
Initial Contact	"CRISPNESS"	(Abrupt . . . . . Mushy)
Surface Rigidity	"HARDNESS"	(Hard . . . . . Spongy)
Final Release	"CLEANESS"	(Clean . . . . . Sticky)

Observations revealed that if the perceptual content of any of these parts was not well represented, the overall percept of a rigid surface contact was degraded and the resulting percept was simply not believable. For example, interaction with a virtual model of a stiff linear spring was found to provide an adequate representation of a *hard* rigid wall when interacting in static contact with the surface, but when interacting with such a model dynamically it was found that the initial dynamic contact had a disturbingly "mushy" or "bouncy" feel which highly distorted the overall illusion of rigidity. A virtual model of a pure linear damper, on the other hand, was found to produce a very *crisp* and abrupt force upon initial contact which can be described as more of a "thud" than a "bounce". Interaction with such a model provided a very realistic sensation of a rigid surface for the first instant of contact, but after that first instant, the illusion of hard contact was lost because the damper lacks static rigidity and allows the joystick to sink into the wall model. When pulling away from a virtual wall modeled as a pure linear damper, the percept again fails because it feels "sticky", as if pulling one's hand out of a thick liquid. This sticky feeling could be eliminated by modeling a virtual damper that only produces an impedance when velocity is towards the wall and has zero impedance when moving away from the wall. Such *directional dampers* were found to provide a *clean* final release from a virtual wall although the initial contact sensation they produced were not as crisp as the pure linear damper. From these first few observations, initial attempts at a perceptual model of a rigid surface contact can be made. A first attempt at a perceptual design of a rigid wall might be to model a boundary layer of intense directional damping to provide the illusion of crisp initial contact and clean final release followed by a stiff linear spring to provide the illusion of a hard static rigidity.

Although perceptual modeling based on personal observation and casual interviews with a few test subjects is insightful, it is the subjective result of exploration rather than objective experimentation. The remainder of this chapter describes an empirical study designed to systematically record the reactions of seven naive test subjects while interacting with various perceptual elements. The goals of this experiment were as follows: 1) to ascertain if the proposed decomposition of the rigid wall contact is a valid and useful way to analyze the percept of a rigid surface contact; 2) to identify which parts of the perceptual decomposition are most important to the overall percept; And 3) to provide insight into what simple

virtual models can be used to provide the salient perceptual features. The intent of this study was to help develop guidelines for the perceptual design of believable rigid surface sensations through the combination of basic perceptual elements. It is also hoped that the methods of perceptual decomposition introduced here will offer designers of virtual percepts an alternative to physical modeling.

4.2. EXPERIMENTAL DESIGN AND METHODOLOGY

4.2.1 Overall Test Design

A two degree of freedom, force reflecting joystick was used to display haptic models of various virtual walls to seven subjects. Subjects were asked to manually explore each of the virtual wall models and use subjective rating scales to quantify the perceptual content of each. For each trial, subjects would interact with a given virtual wall and rate, on scales from 1 to 7, each of the following criteria: *initial surface contact* (crispness), *surface rigidity* (hardness), *final release* (cleaness), and *overall rating* ("wallness"). Subjects were instructed that initial surface contact (crispness) referred to the sensation associated with the very first instant of interaction with the wall model. Surface rigidity (hardness) was described as the sensation associated with applying static pressure to the surface. Final release (cleaness) was described as the sensation associated with the instant of pulling away from the surface. Overall wallness was described as how well the model compared to a real rigid wall contact that was presented to each subject between every trial.

The experiment was divided into two testing sessions, to be referred to as Test I and Test II. Test I exposed each subject to a set of eight simple virtual wall models. Each of the virtual walls used in this test was modeled as a simple element such as a pure linear spring or a pure linear damper. The purpose of using very simple elements in isolation in Test I was to evaluate the perceptual content of basic building blocks from which more realistic percepts could be composed. The primary goals of Test I are enumerated as follows: First, to ascertain if subjects could use the proposed decomposition to quantify the perceptual aspects of given rigid wall sensations; second, to gain insight into how basic elements might contribute to each aspect of the perceptual decomposition; and, third, to correlate relative importance of each part of the perceptual decomposition to the overall wallness ratings.

Test II was run identically to Test I except that subjects were presented with a set of eleven more complex virtual wall models. Each of the virtual walls used in this test was modeled as a combination of two of the simple elements presented in Test I. The primary goals of Test II were, first, to ascertain if subjects could use the proposed decomposition to quantify aspects of more realistic rigid wall sensations;

second, to gain insight into how various combinations of basic building blocks might affect the perceptual decomposition of the overall sensation; and, third, to correlate perceptual decomposition to the overall wallness ratings to understand the relative importance of the three perceptual parts.

4.2.2 Experimental Hardware Setup

Virtual wall models were implemented using one degree of freedom of a high performance two axis force reflecting joystick as previously described in Section 4.1.6. Subjects stood facing the joystick as depicted in Figure 4.6. The joystick handle, which was at a height of 1 m from the floor, was grasped in the right hand. Virtual walls were aligned as shown by the shaded rectangle in Figure 4.6, allowing approximately 7 cm of right to left motion before contact was made. To eliminate spurious haptic information generated by sliding the joystick parallel to the wall surface, motion in the corresponding joystick degree of freedom was blocked with a rigid clamp. To eliminate all audio cues from the testing procedure, subjects were fitted with an headset that presented white noise that masked all sounds from the force reflecting joystick mechanism. To eliminate visual cues from the testing procedure, a partition was used to prevent subjects from viewing their hand as they interacted with the virtual wall models. To keep the four rating criteria fresh in the subject's minds, a sheet of paper was provided with listed the four perceptual criteria and rating scales. A sketch of the testing setup including the joystick mechanism and a test subject can be seen in Figure 4.6.

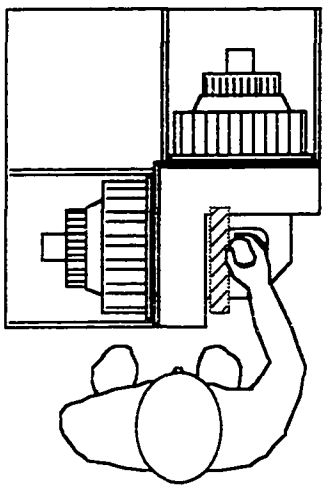

Figure 4.6. Subject standing at joystick. The subject pushes the handle to the left with the right hand to contact the virtual wall (diagonally shaded rectangle). The subject's hand and joystick handle are hidden from the subject's view during experiments.

No restrictions were placed on the type of manual interaction with the virtual wall during the training and data recording sections of the experiments. Subjects were encouraged to challenge vigorously the capabilities of the wall simulations by slamming into the virtual surfaces, withdrawing rapidly, exceeding the resistive force of the wall, and, in general, doing anything to find perceptual flaws in the given model or induce simulation instabilities. They could perform these, or any other, activities as many times as they wished for any particular wall—virtual or real. Furthermore, no time limits were imposed on the overall exploration of any individual presentation of any wall.

4.2.3 Test I Paradigm

Seven Subjects (ages 22-37, all male, all right handed, with no known neuromotor or cognitive impairments) participated in this experiment. The subjects were exposed to a set of eight simple virtual wall models and asked to rate each wall using the four criteria described above: initial contact, surface rigidity, final release, and overall wallness. Between each virtual wall model, subjects were asked to feel a real rigid wall contact implemented by a hard stop placed in the path of the joystick. The complete set of virtual wall models was presented to each subject a total of seven times, each with a random presentation order. The first three passes through the set of eight wall models were used as a training session during which time the subjects were to familiarize themselves with the range of sensations they would be asked to rate. During this training session, subjects were told to concentrate on defining the limits for each rating scale so that they could spread their subjective ratings across the entire range from 1 to 7. Thus, the maximum and minimum limits for each subjective rating scale was defined by each subject individually. Each subject's subjective scales were dependent upon their personal perception of the extreme sensations presented during the training period.

Once the training was complete, data was collected during the following four passes through the set of eight virtual walls. Subjects were asked to rate each of the wall models and verbally report their rating to the experimenter. Each trial consisted of the subject reporting a subjective rating for initial contact, surface rigidity, final release, and overall wallness. During the actual experiment, subjects were asked to concentrate on maintaining consistent rating scales over the four passes through the wall models.

Test I Virtual Wall Test Set

As described above, Test I involved the presentation of a set of eight simple virtual wall models. The wall models used for this test were composed of single elements such as springs or dampers. The

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

complete set of eight virtual wall models studied in Test I along with a graphical representation of the model is shown below in Figure 4.7 for reference.

TEST I VIRTUAL WALL MODELS

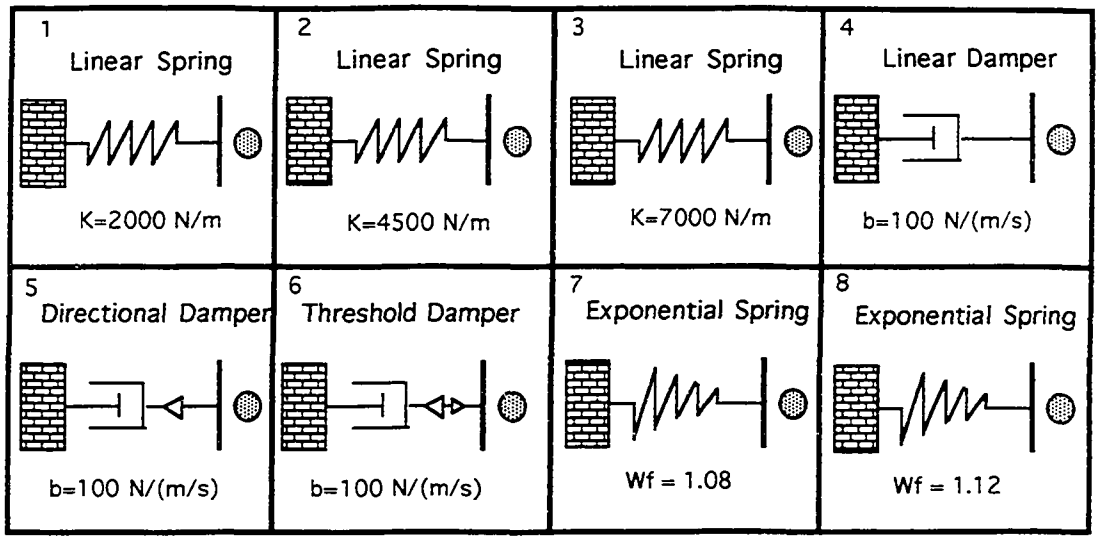

Figure 4.7 Virtual Wall Models Studied in Test I

The following section provides a brief description of each simple element tested: Three of the virtual walls tested were composed of a pure linear spring element with stiffnesses of 2000 N/m, 4500 N/m, and 7000 N/m. These values (which spanned most of the achievable range for the high performance joystick system used) was investigated to gain insight into the effect that increasing stiffness has upon quality of the overall percept. Three of the virtual walls tested were composed of viscous dampers with a damping value of 100 N/(m/s) (approaching the upper limit of what the joystick could produce). One damper model was simply a pure *linear damper* which produced an opposing force proportional to velocity. The second damper, referred to as a *directional damper*, was modeled such that it only opposed velocities in the direction toward the wall and had no effect when moving away from the wall. The third damper, referred to as a *threshold damper*, acted like a typical linear damper that turned into a directional damper when the velocity away from the wall exceeded a small threshold. Directional damper and threshold damper were tested in an attempt to provide damping without producing the "stickiness" associated with heavily damped wall models.

The final two wall models tested were non-linear springs with an *exponential* force Vs displacement profile. The exponential springs were derived to combat a spatial discretization problem inherent to most force reflecting systems which limits the maximum smooth linear stiffness that can be

simulated. This limitation arises in force reflecting systems that employ a servo loop with a digital position feedback from an encoder or A/D converter. In such systems, the maximum linear stiffness of a virtual spring that can be generated is limited by the step size of the smallest digital increment in position feedback. When stiffness becomes large enough that the force increment associated with a single digital step exceeds the user's minimum threshold for perceiving a differential change in force, the sensation is no longer perceived as a smooth linear relation. but rather as discernible incremental steps. The smallest change in force that a person can feel is often referred to by psychophysicists as a *just noticeable difference* (JND). Thus when trying to represent a smooth linear stiffness through force reflecting hardware that uses discrete position feedback, the fidelity of the percept is greatly corrupted when the force increase per digital step exceeds the user's JND for force perception. Even when the joystick is in static contact with a virtual percept, if the force increment per digital step greater than the JND force threshold, a distracting vibration can be felt when the position feedback device bounces on the edge of a discrete step.

In the spirit of design for perception, the question arises: Can a perceptually derived stiffness profile be generated which continually maintains the force increment per-digital-step below the user's ability to perceive a JND in force? Weber [1836] has shown that human ability to perceive a JND is not absolute, but rather varies linearly with the magnitude of the sensation. This linear variation in sensitivity with magnitude is described by a proportionality constant called a Weber Fraction (Wf). Thus, when interacting with a virtual wall, users are most sensitive to small discrete jumps in force when lightly contacting the wall and are significantly less sensitive to such force increments when vigorously interacting with a surface. To take advantage of this linear variation in human sensitivity, a nonlinear spring was derived such that the force increment associated with each encoder pulse always remains below the JND in force. The effect is essentially a spring whose stiffness increases linearly with compression resulting in an overall exponential force versus displacement profile.

At very low force values, the Weber JND proportionality fails to hold. As a result, the initial part of the stiffness profile was derived as a purely linear spring. At the position in the force versus displacement profile where the force increment associated with the linear spring becomes less than the force increment predicted by the Weber proportionality, the profile is transitioned smoothly to the exponential curve proposed above. Although the force profile of this virtual stiffness model is linear for very small values of displacement and exponential for larger displacements, this stiffness profile will be referred to hereinafter simply as the *exponential stiffness profile* to clearly distinguished it from the purely linear stiffness profiles.

Derivation of the Exponential Stiffness Profile

Three parameters completely describe the stiffness profile: the size of a single discrete increment of the position sensor reading *Ds*, the linear force increment *Ft* and the Weber constant *Wf*. For the force reflecting joystick used in these experiments, the size of the optical encoder pulse on the motors dictated the size of the discrete position increment *Ds* to be 0.007 in. The linear force increment is the force increase per discrete position increment to be used in the initial linear region of the stiffness profile. Ideally, this value should be the largest force that remains sub threshold to users of the system. This depends upon the friction and inertia of the system. By trial and error testing, a value of 1 oz was chosen as a conservative estimate for *Ft* for this system. *Wf* is the fraction that describes the exponential part of the profile. Previous studies of manual force resolution suggest that fractions for upper limb force discrimination to be in the neighborhood of 10% [Tan, 1992]. For this experiment, two fractions (Wf=1.08 and Wf=1.12) were chosen for comparison. The variation in fraction was the only difference between the two exponential stiffness models tested.

From these three parameters, the profile can be algebraically computed as follows: *K* is the stiffness of the linear region of the profile and can be computed from *Ds* and *Ft* as follows:

$$K = \frac{Ft}{Ds} \qquad \textit{Linear Stiffness}$$

Knowing the linear stiffness K, the force increment Ft and the fraction Wf, the transition point Xt where the linear increments in force become smaller than the increments described by the fraction JND can be found from the inequality:

$$JND > Ft \qquad \textit{Transition Inequality}$$

Knowing that JND=[(Wf-1) * Net Force] and Net Force = (K Xt) this relation can be rewritten and solved for the transition point Xt:

$$[Wf-1]\ K\,Xt > Ft$$

$$Xt = \frac{Ft}{K\,(Wf-1)}$$

$$Xt = \frac{Ds}{(Wf-1)} \qquad \textit{Transition Point}$$

Thus for all compressions of the spring smaller than Xt, the stiffness profile is linear and is described by Hooke's law as:

$$F = K\,x \qquad\qquad for\ x < Xt$$

For compressions of the spring greater than or equal to Xt, the stiffness profile is described by the exponential relation:

$$F = (K\ Xt)10^{\frac{\log(Wf)}{Ds}(x-Xt)} \qquad for\ x >= Xt$$

Using the relations described above, Test I included two virtual walls modeled as springs with exponential stiffness profiles, both of which were computed with Ft=1 oz and Ds=.007 in. The only difference between the two exponential stiffness models was the fractions used: Wf=1.08 and Wf=1.12. The exponential stiffness curves described by these parameters are shown below in Figure 4.8

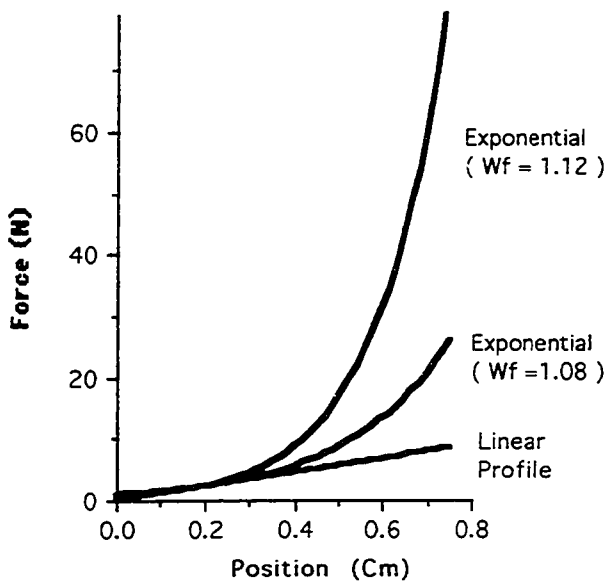

Figure 4.8. Exponential Stiffness Profiles Used as Wall Models

4.2.4 Test II Paradigm

Test II used the same seven subjects who participated in Test I. Test II was run identically to Test I in all ways except in that a different set of virtual wall models was used. While Test I looked at basic elements in isolation, Test II investigated combinations of these basic building blocks into compound virtual wall models. Because the walls presented in Test II evoked significantly more realistic percepts than those tested in Test I, subjects had to develop new subjective rating scales for this test. To eliminate residual effects from Test I upon the subjective scales developed in Test II, subjects were required to take at least a 30 minute break between testing sessions.

In Test II subjects were presented with a set of eleven virtual walls, each of which was modeled as a combination of a stiffness and a damping element investigated in Test I. A basic question of interest was how the ratio of stiffness-to-damping affects the perceptual content of the wall percept. To address this issue, three of the virtual wall models were tested, each of which was composed of a linear spring element in parallel with a linear damper element. The three spring-damper systems were composed such that they had stiffness-to-damping ratios of 800, 120, and 40 sec^{-1}. Also tested were two spring-damper systems implemented as a linear spring element in parallel with a directional or threshold damper element. In addition, two spring-damper models were tested in which the linear damper was contacted by the joystick before the linear spring so as to create a barrier zone of pure damping in front of the standard spring-damper model.

Test II also included systems composed of exponential springs and various damper configurations. One wall was modeled as an exponential stiffness and a damping of 50 N/(m/s). Another was modeled as an exponential stiffness and a damping of 100 N/(m/s). The two remaining walls were modeled as exponential stiffnesses with a directional or threshold dampers of 100 N/(m/s). The complete set of eleven virtual wall models tested in Test II is listed, along with graphical representations, in Figure 4.9.

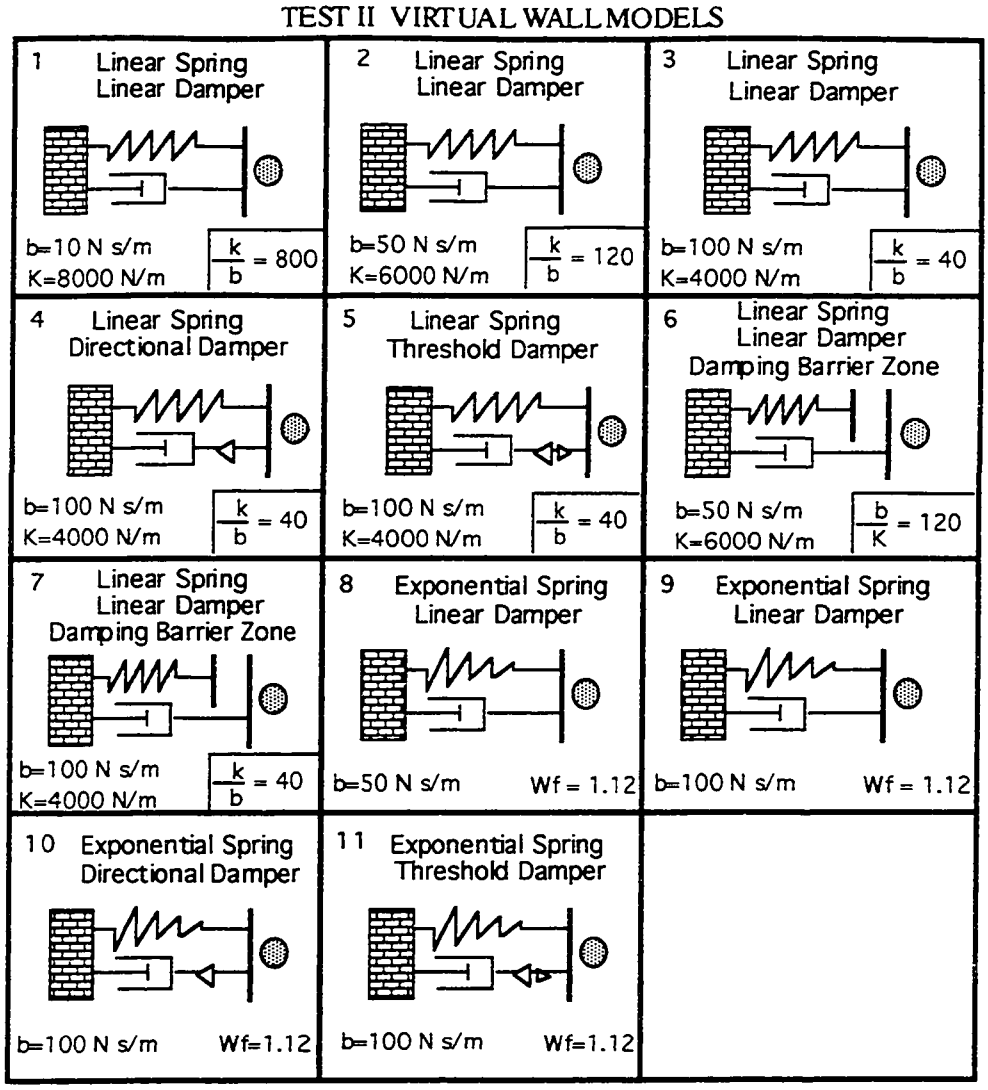

Figure 4.9. Virtual Wall Models Studied in Test II

4.3 DATA PROCESSING

The data collected from the seven subjects during Test I and Test II consisted of integer-number subjective ratings in the range from 1 to 7. For each trial of a particular wall model, independent ratings were recorded for each of the four criteria: initial contact, surface rigidity, final release, and overall wallness.

The actual raw data (consisting of subjective ratings in range from 1 to 7) is inadequate for comparison across subjects because of wide variations among the subjects' rating scales. Each subject developed their own individual rating scales with unique limiting values (i.e. subjects have different impressions of what was meant by the maximum and minimum rating scores of 1 and 7). Each subject also used subjective rating scales with unique non linearities (i.e. the intervals on each subjects subjective scales are not necessarily uniform). In order to compare such rating data across subjects in a meaningful way, the raw data is converted to *rankings* so that non-parametric analysis can be performed as described in Section 4.5. These rankings were computed as the ordinal value of each wall's rating score as compared to the rating scores of the other walls in the experimental set. This was achieved by ordering the wall models from best to worst by subjective rating scores and assigning ordinal values. For cases in which two or more walls shared the same rating values (and thus were indistinguishable by rank) the rank was assigned as the average value. For Test I, which consisted of an experimental set of eight wall models, rankings were assigned such that 1 represented the wall rated the worst of the set and 8 represented the wall rated the best. For Test II, which consisted of an experimental set of eleven wall models, rankings were assigned such that 1 represented the walls rated the worst of the set and 11 represented the walls rated the best. By using ordinal values instead of the raw subjective ratings, the subjects' relative impressions of the haptic percepts were extracted independent of the individual rating scales used. This method of processing allowed for direct and meaningful comparisons of rating data across all seven subjects.

4.4 RESULTS

4.4.1 Raw Data: Subjective Ratings

Before presenting the processed data, the consistency of the raw data should be addressed. The mean coefficients of variation for the raw data (*COV=stdev/mean*) were found to be 0.16 for Test I and 0.24 for Test II. Each coefficient of variation was computed from the four repeated judgments performed by each subject for each wall model using each criterion. These COV values were then averaged over all seven subjects, all wall modes, and all four criteria (NI = 7 x 8 x 4 = 224 ; NII = 7 x 11 x 4 = 308). The COV values reported above reflect relatively low variability in the subjective rating data and suggest that subjects had little trouble decomposing the overall percepts into the given perceptual qualities (initial contact, surface rigidity, and final release). The consistency of the results further suggest that each subject had a clear conceptual model of how each virtual wall compared to the other walls with regard to the perceptual qualities they were asked use. During post testing interviews, subjects reported having little trouble using the given perceptual decomposition and felt that they could clearly distinguish each quality independently of the others. When asked if there were any changes that should be made to the given

decomposition to make the rating process easier, all seven subjects reported being satisfied with the decomposition and none had any suggestions for improvements.

4.4.2 Processed Data: Rankings

As described above in Section 4.3, the raw rating data was normalized by replacing subjective rating scores by ordinal rankings. The complete ranking data for all four criteria over all eight wall models of Test I is shown in graphical form in Figure 4.10 through Figure 4.13. The complete normalized rating data for all four criteria over all eleven wall models of Test II is shown in graphical form in Figure 4.14 through Figure 4.17. For all graphs, the virtual wall models are referenced by the number they were assigned previously in Tables 4.3 and 4.3.

Test I Results

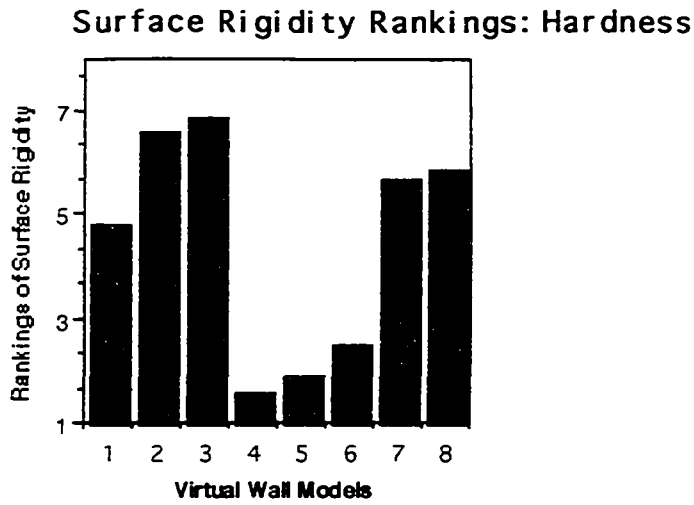

Figure 4.10 Surface Rigidity Rankings for 8 Simple Virtual Wall Models of Test I

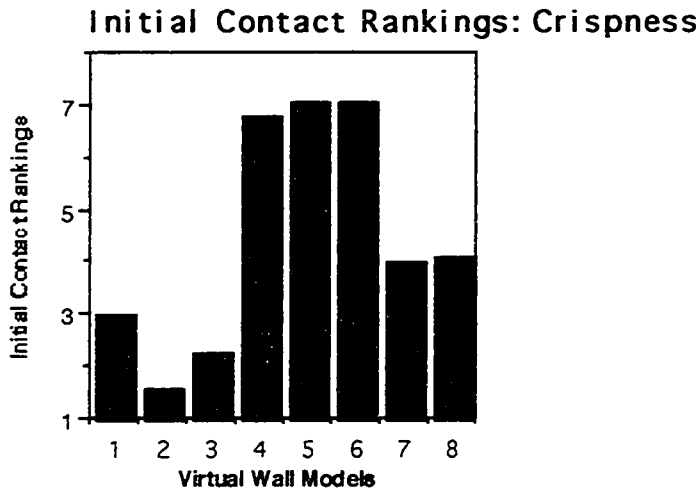

Figure 4.11 Initial Contact Rankings for 8 Simple Virtual Wall Models of Test I

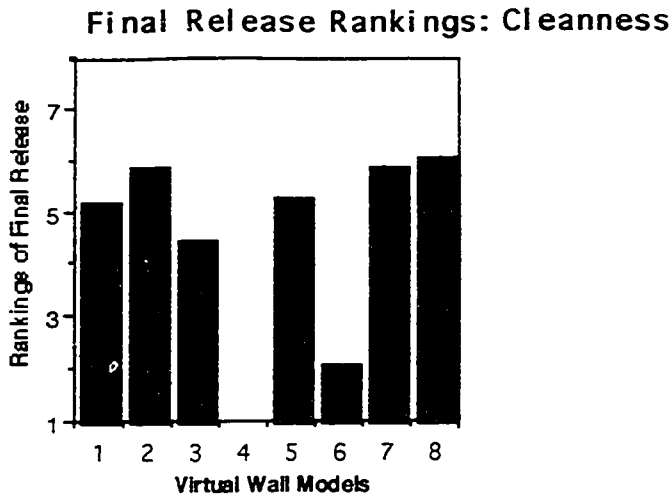

Figure 4.12 Final Release Rankings for 8 Simple Virtual Wall Models of Test I

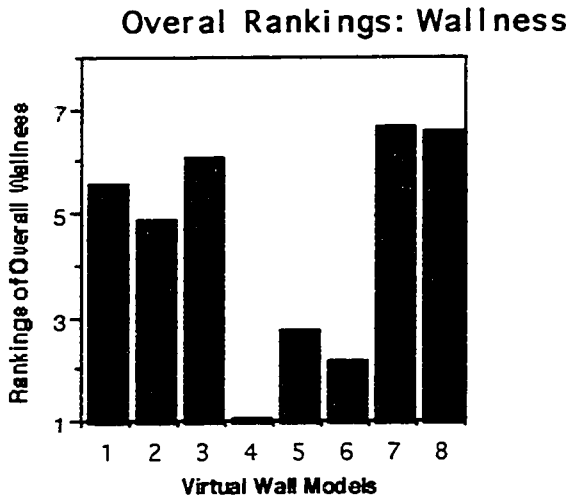

Figure 4.13 Overall Wallness Rankings for 8 Simple Virtual Wall Models of Test I

Test II Results

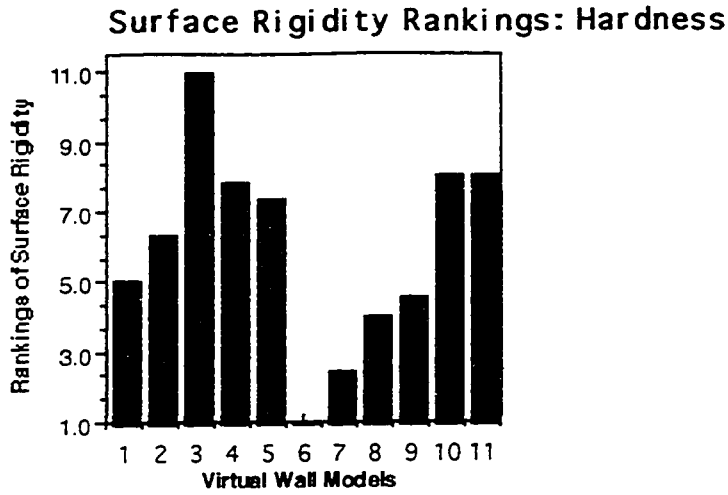

Figure 4.14 Surface Rigidity Rankings for 11 Compound Virtual Wall Models of Test II

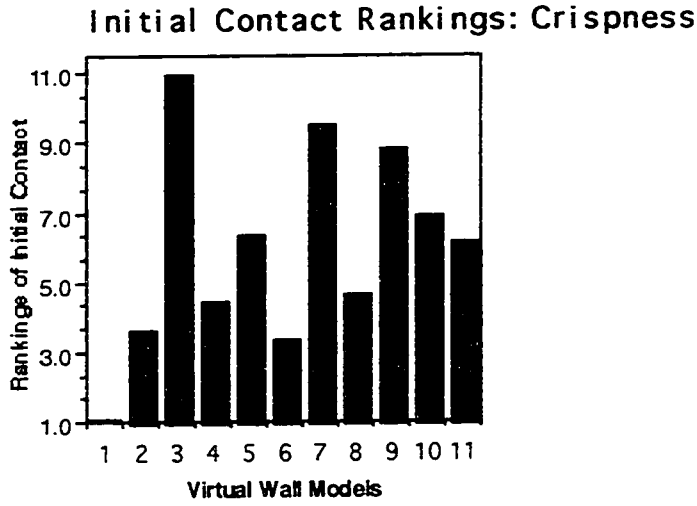

Figure 4.15 Initial Contact Rankings for 11 Compound Virtual Wall Models of Test II

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

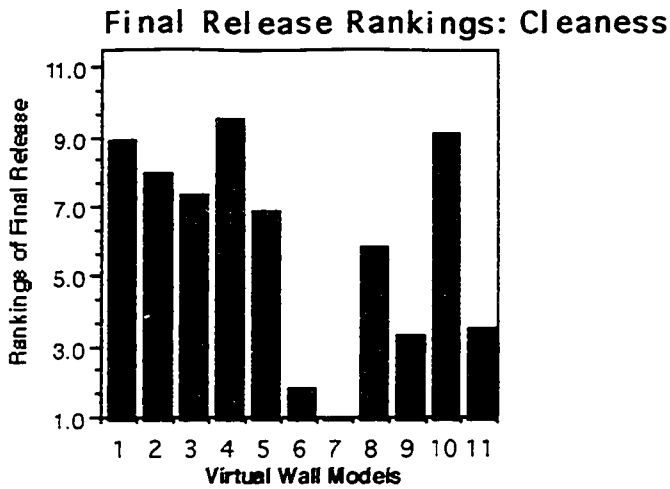

Figure 4.16 Final Release Rankings for 11 Compound Virtual Wall Models of Test II

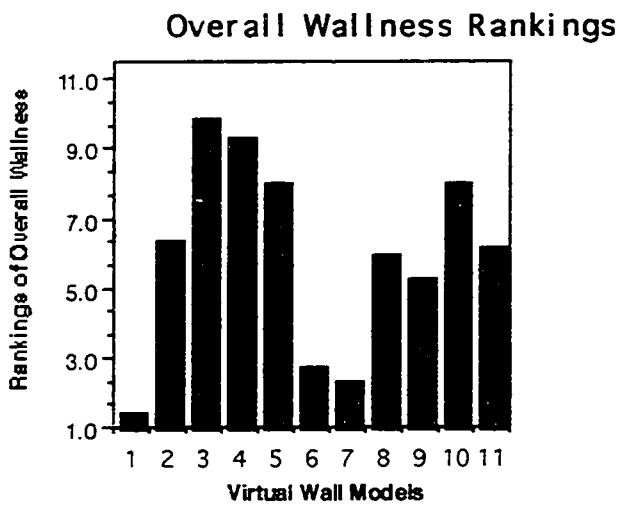

Figure 4.17 Overall Wallness Rankings for 11 Compound Virtual Wall Models of Test II

4.5 ANALYSIS

The following section presents a non-parametric statistical analysis of the processed data presented in Section 4.4. In order to simplify the text, the following abbreviations will be used when referring to the four perceptual criteria addressed in these experiments:

IC = Initial Contact (Crispness)

SR = Surface Rigidity (Hardness)

FR = Final Release (Cleaness)

OW = Overall Ranking (Wallness)

To further simplify the text, the wall models investigated in Test I will be referenced by the number assigned to them in Table 4.2 while the wall models investigated in Test II will be referenced by the number assigned to them in Table 4.3. These are the same reference numbers used in Figure 4.7 and 4.8.

Wall	Model Description	Spring stiffness (N/m)	Damper viscosity (N-s/m)	Fraction
1	Linear Spring	$K = 2000$	—	—
2	Linear Spring	$K = 4500$	—	—
3	Linear Spring	$K = 7000$	—	—
4	Linear Damper	—	$B = 100$	—
5	Directional Damper	—	$B = 100$ for $v < 0$ $B = 0$ for $v \geq 0$	—
6	Threshold Damper	—	$B = 100$ for $v < v_{thresh}$ $B = 0$ for $v \geq v_{thresh}$	—
7	Exponential Spring	—	—	$Wf = 1.08$
8	Exponential Spring	—	—	$Wf = 1.12$

Table 4.2. Test I Virtual Wall Model Summary

Wall	Model Description	Spring stiffness (N/m)	Damper viscosity (N-s/m)	Fraction
1	Linear Spring Linear Damper	$K = 8000$	$B = 10$	—
2	Linear Spring Linear Damper	$K = 6000$	$B = 50$	—
3	Linear Spring Linear Damper	$K = 4000$	$B = 100$	—
4	Linear Damper Directional Damper	$K = 4000$	$B = 100$ for $v < 0$ $B = 0$ for $v \geq 0$	—
5	Linear Spring Threshold Damper	$K = 4000$	$B = 100$ for $v < v_{thresh}$ $B = 0$ for $v \geq v_{thresh}$	—
6	Linear Spring Barrier Linear Damper	$K = 6000$	$B = 50$	—
7	Linear Spring Barrier Linear Damper	$K = 4000$	$B = 100$	—
8	Exponential Spring Linear Damper	—	$B = 50$	$Wf = 1.12$
9	Exponential Spring Linear Damper	—	$B = 100$	$Wf = 1.12$
10	Exponential Spring Directional Damper	—	$B = 100$ for $v < 0$ $B = 0$ for $v \geq 0$	$Wf = 1.12$
11	Exponential Spring Threshold Damper	—	$B = 100$ for $v < v_{thresh}$ $B = 0$ for $v \geq v_{thresh}$	$Wf = 1.12$

Table 4.3. Test II Virtual Wall Model Summary

4.5.1 Test I Analysis

Friedman two way analyses of variance by rank statistics for each criterion (*IC*: $T = 40.75$; *SR*: $T = 38.82$; *FR*: $T = 29.54$; *OW*: $T = 38.67$) allow the null hypothesis—that the subjects show no difference in preference between the wall models of Test 1—to be rejected ($p < .0002$, chi-squared approximation for $df = 7$). This is taken to indicate that the subjects were able to employ the four criteria successfully to distinguish between the eight wall models. Such statistics corroborated post testing interviews in which subjects reported being able to clearly distinguish between the presented walls. The Kendall coefficients of concordance for all four criteria (*IC*: $W = .855$; *SR*: $W = .814$; *FR*: $W = .690$; *OW*: $W = .814$) demonstrate that the seven subjects concurred very strongly ($p < .0001$; chi-square approximation for $df = 7$) in their individual orders of preference for the walls.

Spearman rank correlations to quantify the significance of the relationship between the four criteria (*IC*, *SR*, *FR*, and *OW*) in Test 1 were computed once the rankings of the eight walls by the individual subjects were pooled across all seven subjects ($N = 8 \times 7 = 56$ rankings) and re-ranked. Correlation coefficients for all criterion pairings in Test 1 are listed in Table 4.4

	IC (1)	*SR* (2)	*FR* (3)	*OW* (4)
IC	1			
SR	−.709	1		
FR	−.468	.483	1	
OW	−.578	.787	.615	1

Table 4.4 Spearman Rank Correlations between Initial Contact (1), Surface Rigidity (2), Final Release (3), and Overall Wall (4) rankings for Experiment 1. Note that the correlation matrix is symmetric and that diagonal elements (autocorrelations) are unity.

Each off-diagonal entry in Table 4.4 indicates very significant correlation ($df = 54$; $p < .001$). Thus, not only do the fundamental qualities (*IC*, *SR*, and *FR*) appear to be strongly related with the overall wall (*OW*) percept, but there is also a strong interdependence among the fundamental qualities themselves. The negative entries in the correlation matrix between *IC* and all other qualities confirm that damper wall models, which feel better during the initial contact stage, generally feel worse both in terms of the other two phases of virtual wall interaction (*SR* and *FR*) and in terms of overall wall (*OW*) quality.

While the negative correlation between *IC* and *OW* in Table 4.4 could lead to the unreasonable assumption that good crisp initial contact has a deleterious effect on overall wall quality, partial rank correlation [? 10, 17 ?] when both *SR* and *FR* are held constant reveals that the *IC* phase of wall interaction in Test 1 is *not* related significantly ($r_{S14.23} = .048$, $df = 52$; $p > .50$) to the subjects' judgment of *OW*. With partial rank correlation, however, the associations between *SR* and *OW* when *IC* and *FR* are held constant ($r_{S24.13} = .633$, $df = 52$; $p < .001$), and between *FR* and *OW* when *IC* and *SR* are held constant ($r_{S34.12} = .585$, $df = 52$; $p < .001$), remain highly significant. In other words, the reason that crispness appears to be negatively correlated with overall wallness may be a result of the particular set of wall models tested rather than a perceptual implication of crispness. This effect may have resulted from the simple fact that all wall models in Test I that were high in crispness had to be low in hardness because each model contained *either* a damper *or* a spring element, but not a combination of the two. This artifact will disappear in the following analysis for Test II in which the implemented set of virtual walls allowed for compound combinations of spring-damper elements. Because such combinations were allowed, the perceptual effects of springs and dampers were no longer held mutually exclusive, and the negative correlation between crispness and overall wallness no longer is observed.

4.5.2 Test II Analysis

By Friedman two way analyses of variance for each criterion (*IC*: $T = 54.31$; *SR*: $T = 51.12$; *FR*: $T = 58.94$; *OW*: $T = 50.79$), the null hypothesis that subjects show no difference in preference between the eleven wall models is rejected in Test 2 ($p < 1 \times 10^{-6}$, chi-squared approximation for $df = 10$). Once more this is taken to imply that the subjects were able to employ the four criteria (*IC*, *SR*, *FR*, and *OW*) to distinguish between the given wall models. Also, according to the Kendall coefficients of concordance for each of the four criteria (*IC*: $W = .791$; *SR*: $W = .755$; *FR*: $W = .899$; *OW*: $W = .737$), the seven subjects agreed very strongly ($p < 1 \times 10^{-6}$; chi-square approximation for $df = 10$) with in their rankings of the eleven walls. This suggests a somewhat universal perceptual impression of the various haptic sensations.

Spearman correlation coefficients for all pairings between *IC*, *SR*, *FR*, and *OW* for re-ranked values pooled across all seven subjects ($N = 11 \times 7 = 77$ ranks) are listed in Table 4.5

	IC (1)	*SR* (2)	*FR* (3)	*OW* (4)
IC	1			
SR	.242	1		
FR	–.313	.539	1	
OW	.310	.676	.436	1

Table 4.5. Spearman Rank Correlations between Initial Contact (1), Surface Rigidity (2), Final Release (3), and Overall Wall (4) rankings in Experiment 2. Note that the correlation matrix is symmetric and that diagonal entries (autocorrelations) are unity.

All correlation coefficients in Table 4.5 are statistically significant ($df = 75$; $p < .05$), indicating interdependence among the fundamental criteria (*IC*, *SR*, and *FR*) themselves and between the fundamental criteria and overall wall (*OW*) quality. However, for Test 2, unlike Test 1, partial correlations between each fundamental criterion and overall wall quality remain statistically significant when the other two fundamental criteria are kept constant—*i.e.*, between *IC* and *OW* when *SR* and *FR* are constant ($r_{S14.23} = .320$, $df = 73$; $p < .01$); between *SR* and *OW* when *IC* and *FR* are constant ($r_{S24.13} = .429$, $df = 73$; $p < .001$); and between *FR* and *OW* when *IC* and *SR* are constant ($r_{S34.12} = .317$, $df = 73$; $p < .01$).

4.6 DISCUSSION

4.6.1 Discussion of Rating Results of Test I

The statistical analysis presented in Section 4.5 supports the assumption that subjects can use the proposed perceptual decomposition to distinguish between the various haptic wall models presented in Test I. The following sections independently address each of the four rating criteria which make up the given perceptual decomposition. The discussion focuses on the implications of the ranking graphs depicted in Figures 4.10 through Figure 4.13 as well as the implications of the non-parametric analysis performed in Section 4.5.

Subjective Rating Results of Surface Rigidity (Hardness) for Test I

Turning first to the surface rigidity rankings for the 8 simple virtual wall models investigated in Test I, Figure 4.10 clearly depicts two distinct groupings of hardness rankings. Wall models 1, 2, 3, 7 and 8 are all on the upper end of the hardness ranking scale, with scores ranging between 4.9 and 7.2. The remaining wall models 4, 5 and 6 are all on the low end of the hardness ranking scale, with scores ranging between 1.8 and 2.8. These results demonstrate that all seven subjects had the same impression that virtual walls 1, 2, 3, 7 and 8 provided the strongest sensations of surface hardness while the remaining virtual walls 4, 5 and 6 provided the weakest sensations of surface hardness. The virtual walls 1, 2, 3, 7 and 8 which provided the strong sensations of surface hardness were all modeled as spring elements of various configurations. The virtual wall models 4, 5 and 6 which provided the poor sensations of surface hardness were all modeled as viscous dampers of various configurations.

Of the virtual walls that provided a strong sensation of hardness, wall models 1, 2, 3 were pure linear spring elements with a range of stiffness values: 2000 N/m, 4500 N/m and 7000 N/m respectively. Despite this wide variation in stiffness values, little variation in hardness rankings was recorded among the wall models. While the results show that wall model 1 (2000 N/m stiffness) provided a poor sensation of hardness as compared to the other two linear spring models, there was little difference in hardness rankings between wall models 2 and 3 (4500 N/m and 7000 N/m). Virtual wall models 7 and 8 also provided a well ranked sensation of hardness. These walls were modeled as exponential springs with Weber fractions of 1.08 and 1.12 respectively. The hardness rankings were nearly identical for the two different fractions tested. In addition, little difference in hardness rankings was found between the exponential springs and the linear springs tested. The results suggest the exponential profiles provided a slightly better sensation of hardness than did the 2000 N/m linear stiffness, but were not quite as good as the 4500 N/m and 7000 N/m linear springs.

In direct contrast to the linear and exponential spring models, viscous damper virtual wall models 4, 5 and 6, were found to provide a very weak sensation of hardness. Although virtual wall 4 was a pure linear damper, virtual wall 5 was a directional damper, and virtual wall 6 was a threshold damper, little difference in surface rigidity rankings was noted between these configurations. Of the three, the pure linear damper was ranked worst in hardness sensation and the non-linear threshold damper was the best. The overall results of surface rigidity rankings strongly imply that spring elements of various configurations provide a strong sensation of surface hardness while viscous damper elements of various configurations do not. The results revealed little difference between the five spring model variations tested as well as little difference between the three viscous damper models tested

Subjective Rating Results of Initial Contact (Crispness) for Test I

Turning next to the initial contact rankings for the eight simple virtual wall models investigated in Test I, Figure 4.11 depicts two distinct groupings of initial contact rankings. Wall models 1, 2, 3, 7 and 8 are all on the low end of the initial contact ranking scale, with scores ranging between 1.6 and 4.1. The remaining wall models 4, 5 and 6 are all on the high end of the initial contact ranking scale, with scores ranging between 6.8 and 7.1. These results demonstrate that all subjects had the similar impression that virtual walls 4, 5 and 6 provided the strongest representation of an abrupt and crisp initial contact, while the remaining virtual wall models provided a poor representation of initial contact. The virtual wall models 4, 5 and 6, which provided the strongest sensation of initial contact, were all modeled as viscous dampers. Although model 4 was a linear damper, 5 was a directional damper, and 6 was a threshold damper, no significant difference in initial contact rankings was recorded. The virtual walls 1, 2, 3, 7 and 8 which provided a poor sensation of initial contact were all modeled as spring elements. Models 1, 2 and 3 were pure linear springs of various stiffness (2000 N/m, 4500 N/m, and 7000 N/m, respectively). Of the three, the lowest stiffness value of 2000 N/m was rated the highest in initial contact sensation while little difference was recorded between the 4500 N/m and 7000 N/m stiffnesses. The remaining wall models 7 and 8 were modeled as exponential springs with fractions of 1.08 and 1.12 and were ranked poorly in initial contact. Little difference was recorded in initial contact ranking between the two exponential variations.

The results of initial contact rankings strongly suggest that viscous dampers of various configurations provide a strong sensation of crisp initial contact while various configurations of spring elements do not. This is opposite of what was observed in the surface rigidity ranking trials. This dichotomy between the perceptual representations provided by springs and dampers is confirmed statistically by the negative Spearman Correlation coefficient relating surface rigidity and initial contact as listed in Table 4.4. One explanation of why viscous dampers may provide a strong perceptual representation of

crisp initial contact while spring elements can not, stems from the fact that dampers are dependent upon the derivative of hand position while springs are dependent upon hand position itself. The damper's ability to produce a crisp and abrupt haptic sensation is most likely because the derivative of hand position introduces predictive information into the control of the force stick and thus allows for a faster dynamic response. For example, the incident velocity of the joystick approaching the virtual wall is a predictive indication of how hard the wall is about to be hit. Thus incident velocity provides the control system with lead, allowing an abrupt opposing force to be produced at the instant of manual contact with the wall model. Spring elements have no means of predicting the contact force and thus cannot produce as strong an opposing force until the wall has already been penetrated by some distance. Such penetration preceding the production of a haptic sensation results in a bouncy rather than crisp initial contact feel and can make the wall feel *active* or *live* rather than passive.

Subjective Rating Results of Final Release (Cleanness) for Test I

Turning next to the final release rankings investigated in Test I, Figure 4.12 shows that wall models 1, 2, 3, 5, 7 and 8 are more preferred for final release (cleanness) scale, with scores ranging between 4.5 and 6.1. The remaining wall models 4 and 6 are on the low end of the final release ranking scale, with scores ranging between 1.0 and 2.1. These results demonstrate that the linear spring elements (1, 2 and 3) as well as the exponential spring elements (7 and 8) provided a strong sensation of a clean final release from the virtual wall model. In addition, virtual wall 5 (a directional damper) provided a strong perceptual representation of a clean final release. On the other hand, virtual wall model 6 (a threshold damper) provided a poor perceptual representation of a clean final release. Virtual wall model 4 (a standard viscous damper) provided a less satisfactory representation of final release, being ranked last in cleanness by *all* seven subjects tested. Final release from both of these virtual wall models were often described by subjects in post testing interviews as feeling "sticky". It is important to note that both of these wall models, which were poorly rated for final release, were the highest rated wall models for initial contact crispness. This inverse relation between the perceptual qualities of crispness and cleanness is reflected statistically as a negative Spearman Correlation coefficient relating initial contact (IC) and final release (FR) as shown in Table 4.4.

Subjective Rating Results of Overall Rating (Wallness) for Test I

Finally we turn to the overall "wallness" rankings of the virtual wall models of Test I. As shown in Figure 4.13, virtual walls 1, 2, 3, 7 and 8 were consistently rated near the high end of the ranking scale by all subjects, with scores ranging between 4.9 and 6.7. Virtual walls 4, 5 and 6, on the other hand, were

consistently rated at the low end of the ranking scale, with scores ranging from 1.1 to 2.8. These results demonstrate that all five virtual spring models were given high wallness rankings while all three virtual damper models were given low wallness rankings. This effect is confirmed statistically in Table 4.4 as a strong positive Spearman Correlation coefficient relating surface rigidity (SR) and overall wallness (OW) in addition to a strong negative Spearman Correlation coefficient relating initial contact (IC) and overall wallness (OW).

To gain insight into which perceptual qualities are most important to the overall wallness of a rigid wall percepts, the ranking graph for overall wallness was compared with the graphs generated for surface rigidity, initial contact, and final release. Informal comparison revealed that the general trend observed in the overall wallness graph most closely matched the trend observed in the surface rigidity graph. This result suggests that when simple models are presented in isolation, the most important perceptual quality is the surface rigidity. This observation is supported statistically in Table 4.4 in which the largest positive Spearman Correlation coefficient was the one relating surface rigidity (SR) and overall wallness (OW). Although such a result is insightful, it is still unclear how the relative importance of each perceptual qualities might change when simple elements such as springs and dampers are combined into compound perceptual representations.

4.6.2 Discussion of Ranking Results of Test II

The results of Test I have confirmed that the proposed perceptual decomposition of a virtual rigid wall into *surface rigidity*, *initial contact*, and *final release* provides a viable methodology for dividing a rigid wall sensation into a set of salient perceptual features. The next step is to use such a perceptual decomposition upon a set of more complex virtual wall models with expectations of developing basic guidelines for the perceptual design of more effective rigid wall percepts. Test II attempts to produce such design parameters by performing a perceptual decomposition upon a set of compound virtual wall models composed of parallel combinations of the various spring and damper elements investigated in Test I.

Subjective Rating Results of Surface Rigidity (Hardness) for Test II

Turning first to the perceptual representation of surface rigidity, Figure 4.14 depicts a wide range of surface hardness rankings across the 11 compound virtual wall models investigated in Test II. Wall models 1, 2, and 3 are composed of linear spring and linear damper elements combined in parallel as shown schematically in Figure 4.9 The only difference among these three wall models is the ratio of stiffness-to-

damping. Wall model 1 has high stiffness and low damping, providing a ratio of k/b = 800. Wall model 2 has moderate stiffness and moderate damping, with a ratio of k/b = 120. Wall model 3 has low stiffness and high damping, with a ratio of k/b = 40. Comparing the subjective rankings generated for three wall models, we find that damping ratio has a significant effect upon the perceptual representation of surface hardness. Wall model 1, having the high stiffness-to-damping ratio of k/b = 800, was given an average ranking score of 5.1 which was the lowest of the three. Wall model 2, with a moderate stiffness-to-damping ratio of k/b = 120, was given a higher surface hardness ranking score of 6.4. Wall model 3, with the lowest stiffness-to-damping ratio of k/b = 40, was given the highest ranking of the three that was equal to maximum ranking score of 11.0. The wall with the lowest stiffness-to-damping ratio generated the most convincing surface hardness sensation. This held true for each subject tested.

The fact that these trials revealed an inverse relation between stiffness-to-damping ratio and perceived hardness is particularly interesting in light of the fact that the wall model rated best (with the lowest stiffness-to-damping ratio) had a linear stiffness of K=4000 N/m, while the wall rated worst (with the highest stiffness-to-damping ratio) had a linear stiffness of K=8000 N/m. Thus a wall with significantly lower stiffness but higher damping was unanimously perceived as feeling the *hardest* of the walls tested. Thus the relation between the perception of "hardness" and the physical parameter of "stiffness" is not as simple as one might assume. This is a result that could only have been derived through perceptual testing, being counter-intuitive to what physical modeling would predict. Furthermore it is an empirical demonstration of a fundamental principle on which this thesis is based, namely that physical reality and perceptual reality are distinct realms which should not be treated as one and the same by the designers of virtual percepts. These results suggest that furthur perceptual testing should be done to examine a comprehensive range of stiffness-to-damping ratios and assess their effect upon perceptual qualities.

Wall models 4 and 5 were composed of a linear spring element in parallel combination with a directional damper or threshold damper respectively. Both wall models had a low stiffness-to-damping ratio of k/b = 40. Thus, wall models 4 and 5 were identical to wall model 3 except that both models 4 and 5 were implemented with non-linear damper elements. Despite the addition of the non-linear elements, these wall models showed no advantage in surface rigidity ranking over model 3, With wall models 4 and 5 were ranked lower in surface rigidity than their linear counterpart, wall model 3.

Wall models 6 and 7 included a linear spring with a pure damping barrier zone in front of the spring element, and stiffness-to-damping ratios of k/b = 120 and k/b = 40 respectively. As seen in Figure 4.14, the surface rigidity rankings for these two wall models were the worst of all eleven wall models tested. Wall 6, with the higher stiffness-to-damping ratio, was the worst of the two with a surface

hardness ranking score of 1.1, while wall model 7 was rated only slightly higher with a score of 2.5. This is a particularly interesting result when we compare wall models 6 and 7 to wall models 2 and 3 which differ only in the relative placement of the damper and spring. Although these wall models have identical stiffness and damping parameters, the fact that the damper was encountered a few hundredths of an inch before the spring in wall models 6 and 7 resulted in a dramatic drop in perceived hardness. This results suggests that furthur perceptual testing be done to quantify the effect of spring-damper registration on the perceptual qualities of such a model.

Virtual wall models 8, 9, 10 and 11 were composed of combinations of an exponential spring element and various damper types. Wall model 8 included a moderate linear damper, wall model 9 included a strong linear damper, model 10 a directional damper, and wall model 11 a threshold damper. Wall models 8 and 9 with linear dampers were ranked 4.1 and 4.6 respectively, while wall models 10 and 11 with nonlinear dampers were both ranked 8.1. This result suggests that, the use of directional and threshold dampers in combination with the exponential spring element is preferred over the use of pure linear dampers. This is an interesting result in light of the fact that no advantage was found to using non-linear damper elements in combination with linear spring elements.

The overall results of the surface rigidity rankings of Test II suggest, first, that a damping barrier zone corrupts the representation of surface hardness. The results also suggest that, when implementing a spring-damper virtual wall model (over the range tested), a lower stiffness-to-damping ratio is more important to a hardness sensation than is the net stiffness value used. The results also suggest that, if using exponential spring elements, the use of threshold or directional dampers evokes a better hardness sensation than do pure linear dampers. Overall, the wall model rated best in surface hardness by all subjects was the pure linear spring-damper system with the lowest stiffness-to-damping ratio.

Subjective Rating Results of Initial Contact (Crispness) for Test II

Turning next to the perceptual representation of initial contact, Figure 4.15 depicts a dispersed range of initial contact rankings for the various wall models investigated in Test II. Looking first at the linear spring-damper wall models 1, 2 and 3, we find that damping ratio has a significant effect upon the subjects' perception of contact crispness. Wall model 1, having the highest stiffness-to-damping ratio of $k/b = 800$, was given the lowest possible ranking score of 1.0 by all subjects tested. Wall model 2, with a moderate stiffness-to-damping ratio of $k/b = 120$, was given a moderate ranking score for initial contact crispness of 3.7. Wall model 3, with the lowest stiffness-to-damping ratio of $k/b = 40$, was given the maximum possible ranking score of 11.0 (which implies it was rated best by each and every subject tested).

Virtual wall models 4 and 5 were very similar to the highly rated wall model 3 in that they implement the same damping to stiffness ratio. However, wall models 4 and 5 differed in that they implemented a directional damper and a threshold damper respectively, rather than the pure linear damper used in wall 3. Results displayed in Figure 4.15 demonstrate that these alternative damper configurations do not provide any advantage in producing an effective sensation of initial contact. In fact, the results demonstrate a significant drop in the rankings of initial contact crispness as a result of implementing the non-linear damper elements. The results also demonstrate that the use of a threshold damper provided a mildly better sensation of initial contact than did the directional damper.

Turning to the initial contact ranking results for wall models 6 and 7, which implement a damping barrier zone, we find again that stiffness-to-damping ratio has a dramatic effect upon the perception of initial contact crispness. Wall model 6 implemented a moderate stiffness-to-damping ratio of $k/b = 120$, while wall model 7 implemented a low stiffness-to-damping ratio of $k/b = 40$. The results in Figure 4.15 demonstrate that wall model 6 was ranked 3.4 on average, which was much lower than wall model 7 which was ranked 9.6 on average. The only difference in the models was the lower stiffness-to-damping ratio of wall model 7. In both cases, the wall models with damping barrier zones did not provide a superior initial contact sensation than did the simple linear spring-damper used in wall model 3.

Turning to the initial contact ranking results for wall models 8, 9, 10 and 11 which implement exponential spring elements in parallel combination with various damper configurations, the importance of damping to the crispness of initial contact sensations is again revealed. Wall models 8 and 9 implement exponential springs with linear dampers of 50 N/(m/s) and 100 N/(m/s) respectively. Virtual wall model 9, with the higher damping value, attained a high ranking of 8.4, as compared to wall model 8 which attained a lower ranking of 4.7. Wall models 10 and 11 implemented the same high damping value, but made use of directional and threshold dampers respectively. The addition of these nonlinear damper elements corrupted the initial contact sensation, resulting in slightly lower rankings. Thus a very similar effect was found in the exponential spring trials as was found in the linear spring trials, namely that the highest rated wall model in initial contact crispness had the highest damping value for a given stiffness and made use of pure linear damper rather than the various non-linear elements tested.

The overall results of the initial contact rankings of Test II suggest, first, that low stiffness-to-damping ratio values are of primary importance to the generation of a crisp, believable initial contact sensation. The results further suggest that directional damping, threshold damping, and damping barrier zones do not enhance the initial contact sensation. Finally, the use of exponential springs (as compared to pure linear spring elements) was found not to enhance initial contact sensation of wall models tested .

Subjective Rating Results of Final Release (Cleanness) for Test II

Figure 4.16 depicts a wide range of Final Release rankings across the eleven virtual wall models investigated in Test II. Virtual wall models 1, 2, and 3 consist of linear spring-damper systems of various stiffness-to-damping ratios. As seen in Figure 4.16, wall model 1 with the highest stiffness-to-damping ratio was ranked best of the three in final release. Wall model 3, with the lowest stiffness-to-damping ratio, was ranked worst of the three. Although the variation in final release ranking with stiffness-to-damping ratio is not as pronounced as was found in the surface rigidity and initial contact rankings, it is important to note that the relation revealed here is the inverse of what was found for the previous two perceptual qualities. This conflicting perceptual implication of stiffness-to-damping ratio raises the possibility of a design tradeoff when perceptually constructing an effective rigid wall percept. This tradeoff will be discussed in more detail when we consider the overall "wallness" rankings and address the contributions of each perceptual component to effectiveness of a rigid wall percept.

Virtual wall models 4 and 5 consist of linear springs with directional or threshold dampers respectively. The directional damper was rated best of all walls tested with an average ranking for final release cleanness of 9.6, while the threshold damper was rated significantly lower with an average ranking score of 6.9. Virtual wall models 6 and 7 included a pure damping barrier zone existed in front of the linear spring element. Wall model 6 had a stiffness-to-damping ratio of $k/b = 120$, while wall model 7 implemented a low stiffness-to-damping ratio of $k/b = 40$. Figure 4.16 reveals that both such implementations of damping barrier zones resulted in the worst clean-ness sensations of all walls tested.

Turning finally to virtual wall models 8, 9, 10 and 11 implemented with an exponential spring in combination with various damping configurations, we find the following result: When the exponential spring was combined with a strong linear damping or a strong threshold damping, as used in virtual walls 9 and 11, relatively low cleanness rankings were recorded. When a lighter damping was used as in virtual wall model 8, the final release cleanness ranking was moderately higher. However, we must recall that a light damping is not effective in producing an effective representation of crispness or hardness sensations. When a strong directional damper was used as in wall model 10, the final release ranking score was the second highest score recorded of 9.2. This result suggests that if an exponential stiffness is used in a virtual wall model and high damping is desired to achieve effective representation of crispness and hardness, a directional damper can be used to also achieve an effective representation of a cleanness.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Subjective Rating Results for Overall Wallness for Test II

Considering the overall "wallness" rankings for the eleven virtual wall models of Test II in Figure 4.17, wall model 3 was the most favorably ranked in overall wallness of all wall models tested. It was composed of a linear spring-damper system with a low stiffness-to-damping ratio of k/b = 40. Virtual wall models 1 and 2, composed of linear spring-damper systems with higher values of stiffness-to-damping ratio, were rated significantly lower in overall "wallness". In fact, Model 1, which was implemented with the highest value of stiffness-to-damping ratio, was rated the lowest in overall wallness of all virtual walls tested. These result support the importance of a low stiffness-to-damping ratio in the perceptual design of a believable rigid wall percept.

Virtual walls 4 and 5 were composed of a spring-directional damper or a spring-threshold dampers respectively. These wall models used the same stiffness-to-damping ratio that was used in wall model 3. As shown in Figure 4.17, the replacement of the pure linear damper with a directional or threshold dampers resulted in a moderate drop in the overall believability of the percept. Virtual wall models 6 and 7 implemented a barrier damping zone in front of a linear spring-damper system. These wall models were rated very poorly in overall wallness. This suggests that the spatial registration of spring and damper elements is highly important to the overall perceptual content of a rigid wall model. Furthur testing should be performed to better quantify this registration effect. Virtual wall models 8, 9, 10 and 11 were composed of exponential spring elements in parallel combination with various damper configurations. Figure 4.17 shows that Model 10, composed of an exponential spring and a directional damper, was ranked highest in overall wallness of the four exponential walls tested. In general, the use of exponential springs were not found to provide any advantage over linear spring elements. Nevertheless, future studies of exponential springs profiles should be conducted in which variations in fraction and linear stiffness parameters are fully explored. Since parameters were chosen with little insight prior to this study, optimization may have a profound effect.

Thus far, perceptual testing has revealed that wall model 3, with low stiffness-to-damping ratio, resulted in the most favorable rigid wall percept. The next step is to extend this simple observation into a first pass at general design specifications for the perceptual composition of an effective virtual wall. To develop design parameters for the construction of a rigid wall percept from basic perceptual qualities (crispness of initial contact, hardness of surface rigidity, and cleanness of final release), the contribution of each perceptual quality to the overall effectiveness of wall model 3 should be assessed. Figures 4.14 through 4.16 clearly show that wall model 3 was rated very high in *all three* perceptual qualities. This suggests that all three qualities in the proposed perceptual decomposition contribute significantly to the

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

makeup of an effective rigid wall percept. This result is confirmed in Section 4.5 in which Spearman correlation coefficients are computed to relate each perceptual quality with overall "wallness". As shown in Table 4.5, initial contact, surface rigidity and final release are positively related to overall wallness with coefficients .310, .676, .436 respectively. Thus when designing a virtual rigid wall percept, the most basic design goal should be to represent all three perceptual qualities faithfully. A poor representation of any single perceptual quality has been found to greatly degrade the overall perceptual representation. If we are to provide added importance to any single perceptual quality, results do suggest that surface rigidity, with the highest Spearman coefficient, is the dominant factor. However, we should not take this result to mean that *stiffness* is the most important physical parameter. Results of this study have also shown that the perceptual quality of surface rigidity and the physical quality of stiffness may not necessarily be equivalent. In fact, the best representation of surface rigidity was correlated to a low stiffness-to-damping ratio in this study.

4.7 CONCLUSIONS

Analysis of the raw rating data and rankings from Test I and Test II has provided many insights into the perceptual design of a virtual wall percept. The low coefficient of variation in the subjects' raw scores has suggested that subjects had little trouble in consistently decomposing the overall percepts into the given perceptual qualities. This result supports post testing interviews in which subjects reported having little trouble using the given perceptual decomposition and felt that they could clearly distinguish each quality independently of the others. Statistical analysis in Section 4.5 confirmed that surface rigidity (hardness), initial contact (crispness), and final release (cleanness) are all unique and extractable perceptual qualities of the rigid wall percept although there may be some overlap among them. Future work should ascertain whether additional or alternate perceptual qualities can be devised.

4.7.1 Test I Conclusions

The following section addresses specific results implied by the subjective ranking data for each of the four ranking scales investigated in Test I: *surface rigidity* (hardness), *initial contact* (crispness), *final release* (clean-ness), and *overall ranking* (wallness). The results of surface rigidity testing have demonstrated that various configurations of virtual spring elements were able to provide a convincing sensation of surface hardness while various configurations of viscous damper elements could not. The testing also revealed that the magnitudes of linear stiffness only had a small effect on the overall hardness sensations reported. This result suggests that the hardware performance requirements for the generation of a perceptually adequate representation of surface rigidity could be substantially less than than what physical

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

modeling would predict. It further suggests that perceptual testing be included in the design of virtual percepts to ensure that computational power and hardware design efforts are not being wasted on boosting parameters which may not contribute to the effectiveness of the virtual representation.

The results of initial contact testing have revealed that viscous damper elements of various configurations were able to provide a strong sensation of crisp initial contact while various configurations of linear and nonlinear spring elements could not. Thus, while spring elements in isolation can be used to represent hardness but not crispness, viscous dampers have the opposite effect. This inverse relation between the perceptual implications of springs and dampers strengthen the claim that hardness and crispness are two unique and distinct perceptual aspects of the overall wall percept. It has been hypothesized that the damper's crisp and abrupt feel is due to the fact that it uses the derivative of hand position which provides an element of lead in the control of the force stick. The incident velocity of the joystick approaching the virtual wall is a predictive indication of how hard the wall is about to be hit. Thus heavy damping in the wall model allows the system to respond with an abrupt opposing force at the instant of contact. Wall models with high stiffness but no damping cannot produce a strong opposing force until the wall has been penetrated by the joystick by some amount. This results in a "bouncy" rather than "crisp" initial contact feel.

Results of the final release ranking trials suggest that the linear and exponential spring elements provide convincing final release sensations, while linear damper elements and non-linear threshold damper elements provide poor final release sensations. The non-linear *directional damper* tested was shown to provide as good a final release sensation as the spring elements. Post testing interviews have revealed that subjects describe the final release from a linear dampers and threshold dampers as "sticky", while they described the release from the directional damper or spring elements as "clean". These results, first, confirm that cleanness of final release is an discernible perceptual quality of a rigid wall percept. The results also suggest that if damping is desired in a virtual wall, a convincing final release sensation could be achieved through the use of a directional damping element.

Results of the overall "wallness" rankings of Test I have shown that the most convincing walls tested were exponential or linear spring elements, while the least convincing walls tested were linear and non-linear damper elements. This result suggests that for simple models, the most important perceptual quality of a rigid wall contact percept is the surface rigidity. This was confirmed by a statistical analysis in which the highest Spearman correlation coefficient computed was between surface rigidity and overall wallness. The results further suggest that stiffness values of the linear springs have little measurable effect upon subjects' overall impression of wallness within the range tested (2000 N/m to 7500 N/m). Also,

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

exponential spring elements were not found to provide any measurable advantage in the representation of wallness as compared to linear spring elements.

4.7.2 Test II Conclusions

The surface rigidity rankings recorded in Test II suggest that if a linear spring-damper model is used, a low ratio of stiffness-to-damping is of primary importance a convincing perceptual representation of hardness. In fact, low stiffness-to-damping ratio was found to be more important to the design of an effective hardness percept than was the net stiffness value used. Of the three linear spring-damper walls tested, the wall perceived as *hardest* was actually half as stiff as the wall rated least hard. Such a result could only have been derived through perceptual decomposition, being counter-intuitive to a physical modeling approach.

When using exponential spring elements in parallel combination with various viscous damping elements, results of Test II revealed that the use of threshold or directional dampers evokes a more realistic hardness sensation than do pure linear dampers. The results also suggest that the use of a damping barrier zone in front of a virtual wall model corrupts the illusion of surface hardness and should be avoided. Finally, the results showed that, overall, the wall model rated best in surface hardness by all subjects was the linear spring-damper system with the lowest stiffness-to-damping ratio.

The results of the initial contact rankings of Test II suggest, first, that low stiffness-to-damping ratio is also of primary importance to the generation of a crisp, believable initial contact sensation. The results further suggest that directional damping, threshold damping, and damping barrier zones do not enhance the perceptual representation of initial contact. Finally, the use of exponential springs, as compared to linear springs, was found not to enhance initial contact sensation.

Although low stiffness-to-damping ratio was found to be important for generating hardness and crispness sensations, the opposite was found for the cleanness of final release. For final release, linear spring-damper wall models with the highest stiffness-to-damping ratio were rated best, while wall models with the lowest stiffness-to-damping ratio were rated worst. Because this relation between stiffness-to-damping ratio and the cleanness of final release is the inverse of what was found for surface rigidity and initial contact sensations, it poses a problem for the perceptual design of a virtual percept which satisfies all of the perceptual qualities of a rigid wall used in this study. The results of final release testing, however, do reveal a method of design compromise. The use of a directional damper, in place of a traditional linear damper, provides a convincing final release sensation without corrupting the hardness or crispness of the percept. Thus, a linear spring in parallel with a directional damper provides a convincing representation of

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

all three perceptual qualities. The use of a directional damper was found to have the same beneficial effect when used in parallel combination with exponential spring elements.

The overall wallness rankings support the importance of stiffness-to-damping ratio to the perceptual content of a rigid wall percept. The walls rated best of the eleven walls tested were composed of either a linear spring in parallel with a linear damper or a linear spring in parallel with a directional damper such that the stiffness-to-damping ratio was the minimum value tested. The walls with high stiffness-to-damping ratios were rated the worst in overall wallness of the eleven walls tested. Comparing these results with the rankings of the perceptual qualities (hardness, crispness, and clean-ness), we find the overall wallness to reflect the net effectiveness of the individual perceptual qualities. This suggests that perceptual design of virtual percepts through the combination of basic perceptual elements is a viable means of constructing effective sensations. It should be noted that although stiffness-to-damping ratio was found to be of central importance to the perceptual representation of a rigid wall percept, this should not be taken to imply that designers of virtual walls can construct an effective percept by using the lowest possible ratio of stiffness-to-damping. This preliminary study simply demonstrates that the relative effects of dampers and springs have profound perceptual implications beyond what their physical representations would predict. Additional experimentation is required to fully explore the perceptual design parameters of a rigid wall percept. An effective future study would be one which tests the perceptual implications of a wide range of stiffness and damping values with all other variables held fixed.

The overall wallness rankings also demonstrated that when constructing a virtual wall model, exponential spring elements not *as* effective as linear spring elements. Although no advantage was found for using exponential spring elements over linear models, designers of virtual percepts might consider further testing of exponential spring elements which fully explore the fraction (Wf) and initial stiffness (K) parameters. Because stiffness-to-damping ratio was found to be important to the perceptual representation provided by linear spring-damper models, a similar importance may be attributed to the parameters governing exponential springs.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter Five

"Don't underestimate the power or importance of simple mental aids"

-Donald A. Norman
The Design of Everyday Things

---

5. COMPUTER GENERATED HAPTIC FIXTURES IN 3D

Virtual Fixtures have been defined as abstract perceptual information overlaid on top of the reflected sensory feedback from a telepresence worksites or virtual environments. It has been proposed that virtual fixtures be composed of computer generated perceptual information from a wide variety of sensory modalities including visual, auditory, haptic, and tactile sensations. Studies described in Chapters 2 and 3 have confirmed the effectiveness of virtual fixturing in telemanipulation tasks by demonstrating significant performance improvements resulting from fixture use. Although serving as a confirmation of the potential of virtual fixturing in both direct and time-delayed telemanipulation, initial tests of haptic fixtures were conducted with significant simplifications and limitations which restrict the scope and generality of the findings. The work explored in this chapter addresses the most significant limitations of previous studies (as discussed below) by presenting a new round of subject testing which expands the generality of the virtual fixture concept.

5.1 Limitations of Initial Subject Testing

The most significant limitations of previous results stem from the fact that virtual fixtures were created using *real* physical surfaces to generate haptic sensations rather than using force reflecting hardware to *stimulate* haptic sensations. Although not really "virtual" constructs, such physical implementations of haptic virtual fixtures did provide a powerful testbed to explore the effectiveness of perceptual overlays without having the non-idealities of current state-of-the-art force reflecting hardware influencing the results.

Nevertheless, virtual fixtures are ultimately intended to be used by operators interfaced with force reflecting hardware. Thus, haptic fixturing must be shown to remain effective despite the hardware limitations introduced by force reflecting devices.

In addition to limiting the generality of the results, the use of *real* physical surfaces rather than haptic simulations greatly limited the variety and composition of haptic percepts that could be presented to subjects. As a result, the only haptic fixtures implemented in initial tests were composed of simple combinations of rigid planar surfaces. Although such tests did demonstrate the effectiveness of virtual fixtures and supported the premise that fixtures could be composed of simple building blocks, the testing of *only* rigid surfaces provided very limited insight into the potential of haptic fixturing. Also, the testing of only hard surfaces gives the limited impression that haptic fixtures are effective primarily as a means of providing rigid kinematic constraints upon the users kinetic activities. Thus, it is important to the general exploration of the fixturing concept to test haptic sensations other than simple rigid surfaces.

Another limitation resulting from the use of real physical surfaces rather than haptic simulations, stems from the fact that the physical surfaces are restricted by *physical* design parameters such as clearance, robustness, occlusion of the workspace, and other mechanical constraints. As a result of these design restrictions, the initial studies in Chapters 2 and 3 were limited to the implementation of *two-dimensional* virtual fixtures. Although the use of two-dimensional fixturing provided an effective platform for demonstrating the virtual fixture concept, virtual fixtures for general purpose telemanipulation in three-dimensional workspaces requires that fixtures be implemented in three-dimensional space. Thus to provide a more general confirmation of the virtual fixture concept, three dimensional virtual fixtures should be explored.

5.2 Follow-up Study to Expand Generality of Virtual Fixture Concept

In order to provide a more general investigation of virtual fixturing, a follow up study has been conducted which more closely reflects the intended application and implementation of virtual fixturing. This study generates haptic virtual fixtures as computer simulated constructs by using a full upper-body force-reflecting exoskeleton to display haptic sensations to the user. Since previous tests in Chapters 2 and 3 used physical elements to represent haptic fixtures, the current study served as the first empirical exploration of virtual fixtures in which the haptic percepts were truly *virtual*. Because this study uses a force reflecting exoskeleton to generate haptic sensations, the virtual fixtures presented in this chapter are not limited to rigid surfaces as were previous studies. Thus this study was the first exploration of the virtual fixturing in which abstract haptic percepts could be implemented. In addition, the use of the force

reflecting exoskeleton allowed for the generation of three-dimensional haptic models. Thus this study was the first exploration of virtual fixturing in which three-dimensional percepts could be displayed.

5.3 EXPERIMENTAL HARDWARE

A telemanipulation peg-insertion paradigm was again chosen for this round of empirical testing. The basic goal of the study was to quantify the effectiveness of computer generated virtual fixtures in enhancing operator performance in a telepresence task. The peg-insertion task was used because parameters can be tightly controlled and human performance can be readily quantified. The following sections describe the extensive hardware setup and experimental protocol which made this round of empirical testing technically possible.

5.3.1 Overall System Setup

These tests required an extensive array of hardware and software to be designed and assembled into a functional system. A large number of sub-systems had to function in unison to allow human operators to naturally control a robotic manipulator in a remote telepresence worksite while at the same allowing haptic virtual fixtures to be overlaid on top of the user's perception of that worksite. Two technical challenges dominated the design of the experimental setup. The first challenge was to create an effective telepresence link between the operator and the remote manipulator. As postulated by Loomis [1993], an effective sense of presence can only be established in a remote environment when a consistent and predictable relation exists between efference and afference. Thus the experimental hardware had to allow for a natural mapping between the users kinesthetic activities (efference) and the motions of the remote manipulator. Similarly, the hardware had to provide a means of reflecting such kinesthetic activities to the user through appropriate sensory feedback from the remote worksite (afference). The second technical challenge was to allow for three dimensional haptic virtual fixtures to be generated and presented to human subjects such that they are perceived as being *overlaid* on top of the real remote environment. This required not only integration of hardware, but also required a formal calibration procedure for assuring the virtual and real environments were in registration.

To meet the technical challenges described above, the resulting experimental system involved many computationally intensive operations to occur simultaneously such as tracking the operators motions, modeling the virtual environment, controlling the remote manipulator, conducting the experimental trials, collecting data, as well as monitoring the safety of the human subject in the force reflecting hardware. As a result of such computational complexity, the resulting hardware system is a collection of networked subsystems , each having its own microprocessor and drive software to share the burden. Although the

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

development of this system was only a byproduct of this thesis. I will take some time to describe the setup simply because the end result is one of few telepresence environments ever constructed with haptic display.

As shown in Figure 5.1, the overall experimental setup can be divided into two primary subsystems: the *operator space* and the *remote environment*. The operator space is where the human subject is interfaced to the system and consists of the FREFLEX upper body exoskeleton (manufactured by Odetics Inc. under contract by USAF), the stereo vision system, and two microprocessors on a VME bus which share the processing requirements of controlling the exoskeleton. The remote worksite consists of the MERLIN robot arm, the robot controller, the peg insertion Task Board, and two microprocessors which share the load of running the experiment and handling communication between the exoskeleton and the robot arm. The following sub-sections describes each of these subsystems as well as the communication pathways which allow this distributed system to operate in unison.

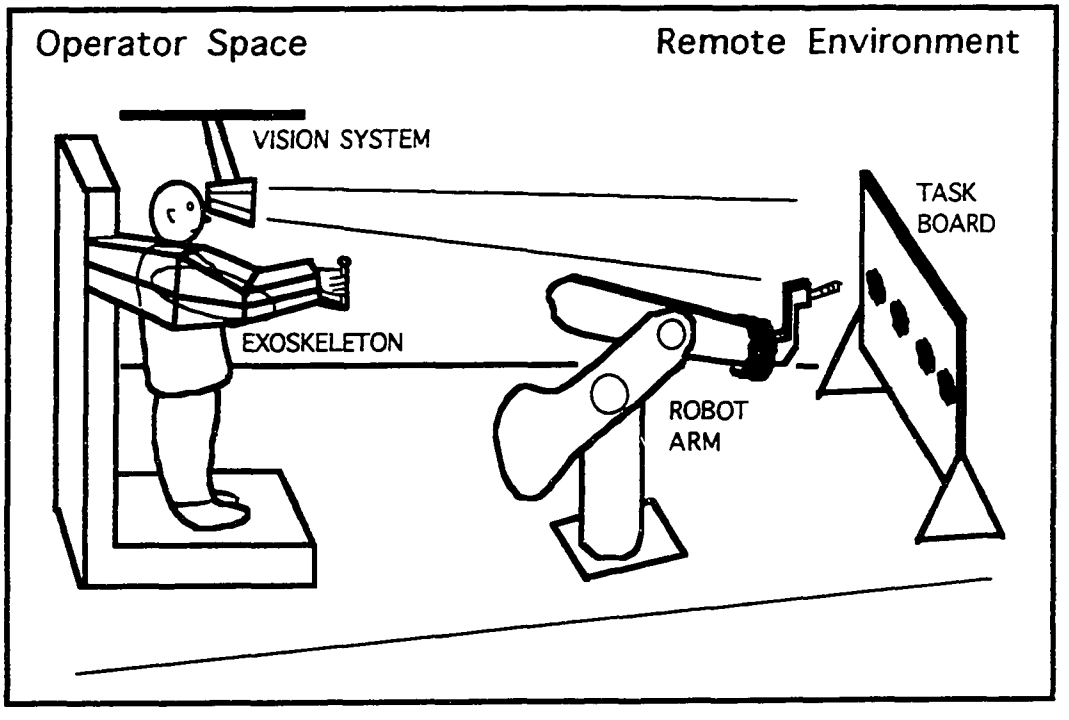

Figure 5.1: Overview of Experimental Hardware Setup

5.3.2 The FREFLEX Force Reflecting Exoskeleton

The FREFLEX upper body exoskeleton is a seven degree of freedom, anthropomorphicly mapped force-reflecting master designed and built by Odedics inc. The system is worn by a user by sliding his/her

right arm into the exoskeleton while standing on the control platform. The kinematic structure of the exoskeleton was designed to allow a natural range of arm and wrist motion with minimal constraint. Software controllers perform gravity compensation and friction compensation to greatly reduce the encumbrance imposed by the device on the user. It should be noted that the current system does not provide dynamic compensation so although the system feels weightless, the full inertia of the exoskeleton is felt by the user.

Because this system is essentially a powerful robot with a person strapped inside, it includes a wide variety of safety features to protect the user from injury. These precautions include hard-wired velocity shut down, a dead-man hand switch, and a safety foot pedal, all of which all have to be engaged for the system to operate. Every step possible was taken to assure the safety of the human subjects placed within this system. For this set of experiments, hardware and software limits were set to restrict the maximum force output at the wrist to be no greater than 15 lbs. This value was chosen because it was strong enough to produce rich haptic sensations but weak enough that it could be overcome by the subjects if a problem arose.

During operation, the angles of all seven rotary joints on the exoskeleton structure are measured by inline potentiometers and fed through the forward kinematic equations to achieve the Cartesian position and orientation of the hand grip with respect to the users reference frame. These positions and orientations are computed by a 68030 microprocessor on a VME BUS and placed in shared memory so it can be quickly accessed by the micro controller which commands the MERLIN robot arm. Such rapid information transfer between the exoskeleton and slave robot allows for efficient kinetic tracking in the telepresence interface. In addition to sending out a position vector, the FREFLEX controller receives as input a Cartesian force vector representing the magnitude and direction of the force to be applied to the user at the hand grip. This vector is multiplied by a decoupling matrix to compensate for coupling effects in the cable transmission and then multiplied by the inverse Jacobian to achieve joint torques. These joint torques are then software limited to be within safe magnitudes and then fed to each joint's motor controller.

All the processing for exoskeleton control is shared by two microprocessors, a Motorola 68030 and an Intel 486. Despite the distributed architecture, the bandwidth of the system is greatly limited by the large amount of computation required each cycle. For this set of experiments, the update rate of the FREFLEX controller was approximately 300 Hz. This value has been quoted by Brooks [1990] as the lower acceptable bandwidth limit for force reflection to the human hand. Although slow for a haptic display, informal observation found this update rate acceptable for the production of haptic percepts. Although it was not possible to produce a very crisp rigid surface with such a slow update rate, the purpose of this set of experiments was to confirm that virtual fixtures remain effective despite distortions introduced

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

by hardware limitations. Thus the relatively low bandwidth of this system actually made for a suitable platform for this set of tests.

5.3.3 The Task Board (Air Force/Navy Teleoperator Performance Evaluation Battery)

While the same task-board was used in this round of peg-insertion testing as was used in Chapters 2 and 3, a different arrangement of holes was employed in the current study (see section 2.2 for complete description of task board hardware.) This time, four 2.00 cm diameter holes were used, spaced to allow horizontal, diagonal, and vertical peg movements of approximately 8.0 cm, 5.7 cm, and 8.0 cm respectively. As shown in Figure 5.3, these motions were chosen as a linearly independent spanning set of all linear motions in the vertical plane.

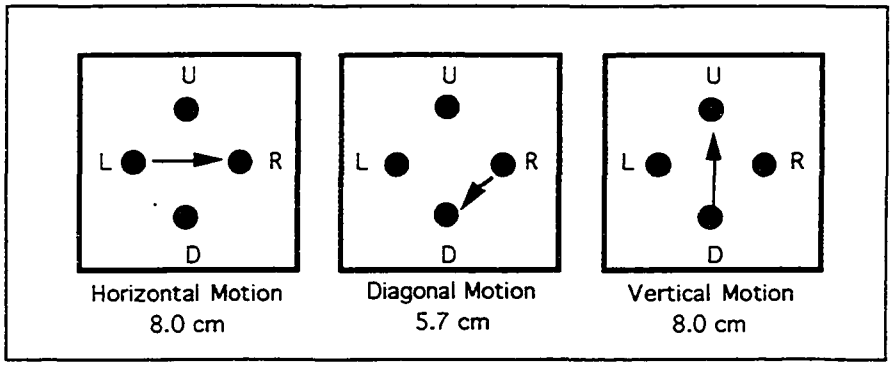

Figure 5.3: Horizontal, Diagonal, and Vertical Peg Motions Tested

5.3.4 The Merlin Robot Arm:

For these experiments, the position of the subject's right hand was tracked by the FREFLEX exoskeleton and fed to the Merlin robot controller with a 1:1 scaling. Thus three dimensional motions of the subject's hand were naturally mapped to command the robot's end effector. A special end effector was designed for the Merlin robot which housed the peg and was equipped with a single axis force transducer composed of a spring and a linear position sensor. With this transducer, axial force into the task board could be measured and reflected back to the FREFLEX exoskeleton. Thus, when the slave robot moved the peg into contact with the task board, the user would feel the contact through the force reflecting exoskeleton. Although tangential forces were not reflected, the normal force feedback was a significant improvement over previous studies in which no force feedback was reflected from the remote environment. Because only axial force feedback was provided to the user, the slave robot was programmed to maintain its end-effector orientation perpendicular to the board at all times. This arrangement made tangential forces

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

insignificant for task execution as well as provided system robustness by eliminating the possibility of subjects torquing the peg in the hole and damaging the setup.

5.3.5 The Stereo Vision System

Although the previous studies of virtual fixtures used a monocular vision system which proved effective for the peg-insertion task, complaints of user discomfort during these initial tests lead to the design of a stereo vision system for this new round of experiments. In order to achieve high resolution visual feedback from the remote environment, a simple but effective stereo vision system was designed. The system was developed as an inexpensive means of creating the visual illusion of operator presence in the remote workspace. The system used 7X power binocular optics focused upon the distant remote worksite with a perspective consistent with that which a user would see if standing directly before the task board. In order to achieve the proper visual perspective and magnification while preventing the robot arm from obscuring the user's view, the FREFLEX system was placed on a raised platform 2.5 feet off the ground. The stereo vision system was fixed to the ceiling of the room on an adjustable support which allowed the height and angle to be set appropriately for each subject.

When looking through the vision system, a user saw the task board as if it were situated just two feet in front of them. The user could not see their own hand when looking into the vision system, but rather saw the robot end effector in the same location which their kinesthetic sense suggested their hand should be. This correlation between kinesthetic and visual feedback from the remote site provided a primary mode through which afference matched efference and thus strengthened the user's sense of presence within the workspace. Post testing interviews confirmed that all subjects were comfortable with the vision system and felt that the visual feedback provided them with a strong sensation of presence in the remote worksite.

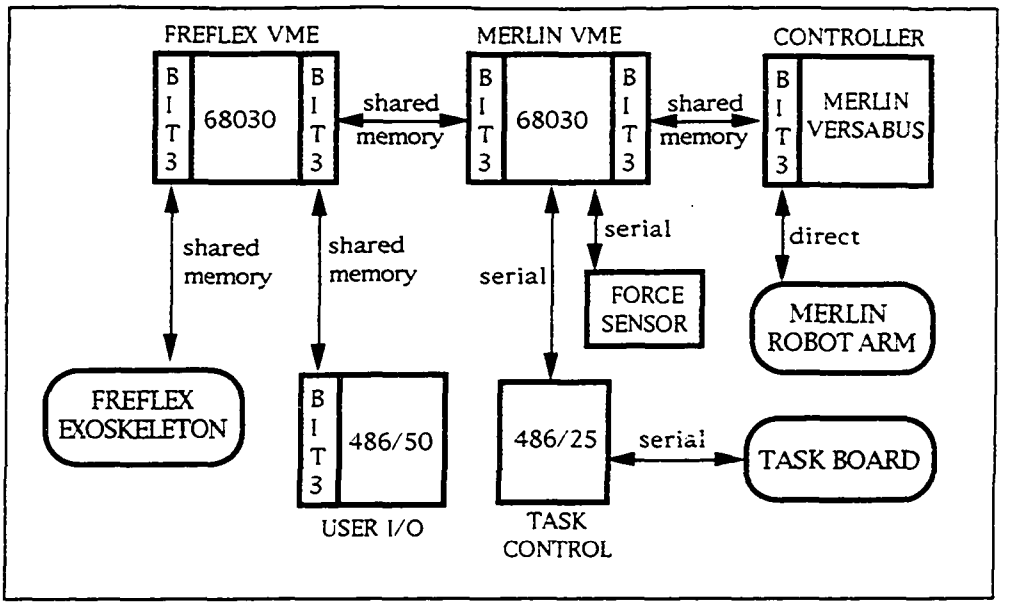

Figure 5.4: Computer Architecture of Experimental Hardware

5.3.6 Computational Architecture

As described for each subsystem, many microprocessors were used in the overall experimental system which functioned in parallel and needed to communicate continually during system operation. Parameters such as operator hand position which required rapid communication rates were handled using shared memory on a VME-BIT3 card. Parameters such as the state of the task board which did not require high bandwidth were transmitted over RS232 serial lines running at 115Kb. Figure 5.4 depicts the overall distributed architecture developed to run this experiment.

5.4 EXPERIMENTAL DESIGN AND METHODOLOGY

A group of 7 human subjects were put through an extensive battery of tests designed to evaluate the effectiveness of virtual fixtures in a telemanipulation task. The overall procedure, which included extensive training and data collection trials, required the daily testing of each subject over a three week period. A peg-insertion task was used as the basic paradigm for subject testing. The overall testing procedure was broken up into three parts to be referred to as Test I, Test II, and Test III. Test I was a baseline test of basic human performance in which subjects performed the peg-insertion task *in person* by standing at the task board. Test II was a baseline test of telepresence performance in which subjects performed the peg-insertion task through the exoskeleton system but were not provided with virtual

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

fixturing. Test III was the comparison test in which subjects performed the peg-insertion task through the exoskeleton system and were required to use each of the five virtual fixture designs.

5.4.1 The Peg Insertion Task

The basic task required of subjects was to remove a peg from a designated start-hole and insert the peg in a designated target-hole as quickly as possible. Data collected from the trials consisted of the time elapsed from the instant the peg left the start-hole until it was fully inserted in the target-hole. Although previous testing of two dimensional virtual fixtures used a formal Fitts' Law paradigm, the new round of testing used a modified peg-insertion paradigm to reduce the number of trials required for testing three dimensional virtual fixtures. This modified Fitts' Law paradigm reduced the trial count by using only one peg size (1.50 cm) rather than the set of three peg sizes used in previous studies. Although this study did not test a range of task difficulties as did previous tests, this study tested three different peg motions: horizontal, diagonal, and vertical. It was felt that testing a set of peg motions which spanned the task plane was a more effective way to explore three dimensional fixture designs than simply testing multiple task difficulties.

Four peg holes of 2.0 cm diameter were used in this study which were equally spaced 4.0 cm from the center of the task board as shown in Figure 5.5. These holes were referred to as Left, Right, Up and Down, indicated on the figure by L, R, U, and D respectively. The holes were closely spaced so as to only require peg motions of small amplitude. Small amplitude motions were chosen for this study to reduce the encumbrance imposed by the added inertia of exoskeleton. It was found that small amplitude motions could be performed for many trials without significant fatigue. Since a small amplitude motion was used, a tight tolerance between peg and hole was used to maintain a high task difficulty level. The 1.5 cm peg used in this experiment made for very challenging insertion when performed remotely though the telepresence system. In fact, three 30-minute practice sessions were required before all subjects became proficient at this task. Subjects were considered proficient when their coefficient of variation across all trials of a single session was below 30%.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

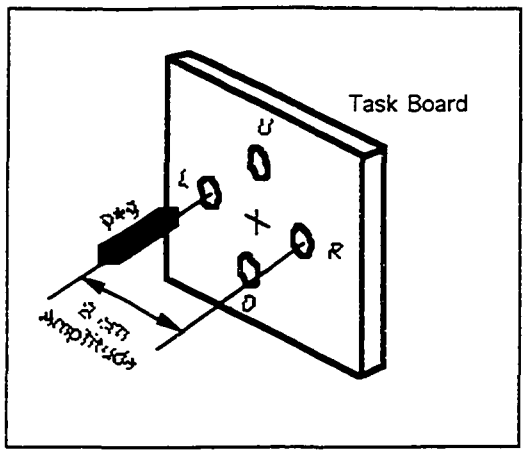

Figure 5.5: Peg Insertion Task

5.4.2 Testing Procedure:

Each of these tests, (Test I, Test II, and Test II), were broken up into a number of daily sessions which included a short warm-up followed by 30 recorded trials. Each 30-trial data collection session was composed of 10 peg-insertion trials for each of the three peg motions described above (horizontal, diagonal, and vertical). Each testing session only lasted between 15 to 25 minutes so that fatigue would not greatly effect the results over the course of the session. The testing procedure required subjects to perform each 30 trial session at their own pace by executing a pre-defined sequence of peg insertion motions- (first horizontal, then diagonal, then vertical, then back to horizontal, etc.). By using this interlaced sequence of different peg insertions tasks, subjects could not easily memorize the physical motions required of the task. This made for a robust and consistent testing procedure in which the performance of fully trained subjects did not vary much over the course of the data collection sessions. Also, by using a pre-defined sequence of peg insertion motions, subjects could perform the entire testing session with little (if any) interference or distraction from the experimenter. This was important in preserving the subject's concentration and maintaining the illusion of presence in the remote worksite. Furthermore, by allowing subjects to perform the task at their own pace, subjects were given the opportunity to rest at any point if they felt that they were getting fatigued. The overall testing schedule is tabularized below in Table 5.1.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Table 5.1: Overall Testing Schedule

TEST	SESSION	DESCRIPTION OF TRIALS
Test I		In Person Peg Insertion Trials
	Session 1	practice trials
	Session 2	practice trials
	Session 3	practice trials
	Session 4	data trials
Test II		Telepresence Peg Insertion Trials
	Session 1	hardware familiarization
	Session 2	practice trials
	Session 3	practice trials
	Session 4	practice trials
	Session 5	data trials
Test III		Telepresence Peg Insertion Trials with Virtual Fixtures
*	Session 1	Intersecting Surface Fixtures
*	Session 2	Uniform Damping Fixtures
*	Session 3	Snap-To-Plane Fixtures
*	Session 4	Variable Damping Fixtures
*	Session 5	Snap-To-Line Fixtures
	Session 6	no fixture baseline

* denotes that order of these trials were randomized

5.4.2.1 TEST I: In-Person Peg Insertion Trials

The intent of the in-person Test I trials was to record base-line readings for the peg-insertion task free from performance degradation resulting from remote presence. The ultimate goal was to compare the performance achieved in a telepresence task (with and without virtual fixtures) to the performance achieved by a person performing the same task "in-person." Thus it was necessary to create an "in-person" task identical to the remote presence in every way possible, except for the single distinction that the task was performed in-person rather than through a telepresence interface. This posed interesting challenges because the emcumberances imposed by the exoskeleton during telepresence activity was not imparted during the "in-person" peg-insertion task.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

One encumbrance is due to the fact that the operator interaction with the exoskeleton occurs through a palm-grip - i.e., the operator of the telepresence system could not make use of finger dexterity when performing the peg insertion task. An equivalent in-person task should not allow users to hold the peg between their fingers and thus make use of finger dexterity when inserting pegs. Another encumbrance was that the exoskeleton only reflected axial (out of the board) forces to the user, but not lateral forces or torque's on the peg. An equivalent in-person task should only allow users to feel axial forces out of the board. Finally, although the exoskeleton did compensate for its own weight, it did not perform dynamic compensation, thus imparting significant inertial effects upon the operator during the task. An equivalent in-person task should therefore impose equivalent inertial effects upon the operator but not impart any gravitational loading.

In order to compensate for these effects and make the in-person and remote presence cases as similar as possible, a human interface tool called the *Exoskeleton Encumbrance Simulator* was designed for this experiment. This tool, as shown in Figure 5.6, was designed to simulate the encumbrances imposed by the exoskeleton in the following ways. First, the peg was fixed to a palm-grip so that the operator interacted with the peg using the same muscle groups as when using the exoskeleton. Second, the peg was mounted on a cantilever spring which was rigid for axial loads but somewhat flexible for lateral loads and torques. Thus the user primarily felt axial forces during interaction with the task board as was the case when using the exoskeleton. Finally, the hand grip was fitted with lead ballast which simulated the mass of the exoskeleton arm. Since the exoskeleton performed its own gravity compensation, this ballast was counter-weighted with lead blocks on the opposite end of a pulley system. Thus the hand grip simulated the increased inertia of the exoskeleton while compensating for its own weight as did the exoskeleton.

In addition to the hand grip, the task board was set up for Test I so that the operator would assume a very similar stance when performing the task in-person as they would when using the exoskeleton. The task board was set at a height such that subjects would have a very similar view of the task board when performing the task in-person as they would get when using the stereo vision system in the remote workspace. Although this system did not result in an in-person task which was identical to the remote presence task, it did provided a close enough approximation to be very useful for this experiment.

The experimental procedure for Test I went as follows: Each subject was introduced to the task board and peg-insertion task. They were told that their goal was to perform the peg-insertion as quickly as possible (as long as they could consistently and confidently perform the task at that speed). Next they were asked to perform three practice sessions (one each day for three consecutive days). Each practice session began with 12 warm-up trials to allow subjects to get used to the feel of the hand grip and the peg-insertion task. The warm-up was followed by 30 timed trials. At the end of the third practice session, results

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

confirmed that learning was no longer having a significant effect upon performance. On the forth day, subjects were asked to perform a final test session which was used to collect the actual data for this experiment.

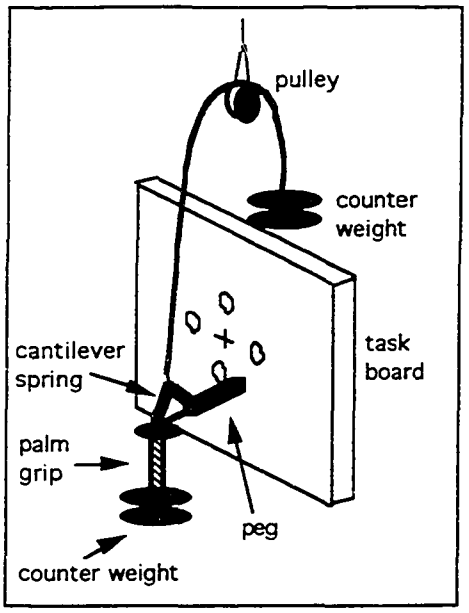

Figure 5.6: The Exoskeleton Encumbrance Simulator shown with task board

5.4.2.2 TEST II: Telepresence Peg Insertion Trials without Virtual Fixtures

After all subjects had gone through Test I, they were introduced to the experimental hardware to be used in the telepresence peg-insertion trials. Subjects were placed in the operator space which included the FREFLEX exoskeleton and the stereo vision system. Subjects were then taught how to operate the safety features for the exoskeleton system which included a hand switch and foot pedal which engaged the force reflection motors. Subjects were then introduced to the remote workspace which included the MERLIN robot arm and the task-board. All subjects were given a basic understanding of how the system operated, how forces were reflected, and how commands were transmitted to the remote site.

Subjects were provided with an initial 30 minute session during which time they were asked to familiarize themselves with the equipment. Many subjects needed this exposure period to simply get aquaitinted with the illusion of telepresence, to get used to the feel of the exoskeleton, and to get

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

comfortable with controlling the remote manipulator. Once comfortable with the system, all subjects were asked to perform three practice sessions (one each day for three consecutive days) during which time they would become fully trained on the telemanipulation task. Once training was complete, subjects were asked to perform a final testing session which would serve as the data collection period.

5.4.2.3 TEST III: Telepresence Peg Insertion Trials with Virtual Fixtures

Once the subjects had gone through Test I and Test II, they were then introduced to the notion of virtual fixturing. They were told that the exoskeleton would overlay force information on top of their perception of the remote environment and that these perceptual overlays were called virtual fixtures. Subjects were told that a set of five different virtual fixture designs would be implemented and that these fixture were intended to enhance their performance in the peg-insertion task. Although all subjects were asked to make an attempt at using these virtual fixtures to enhance their performance, no guidance was given as to how such fixtures should be used. A five day testing period was implemented in which subjects would perform a short warm up followed by 30 data collection trials during each testing day. A new virtual fixture design was introduced each of the five days. The order in which subjects were tested on the five different fixture designs was randomize to avoid ordering effects.

Unlike the previous two tests, Test III did not provide subjects with an extensive practice period which otherwise would guarantee that they would be fully trained on the given task. In fact, subject were only allowed a short 12 trial warm-up session with each virtual fixture design before performing the data collection trials. There were a number of reasons for not training subjects on the various fixture designs. First, all subjects had already been fully trained on the peg-insertion task (without fixtures) and were comfortable with using the hardware to perform peg-insertions. Thus there was less of a need for an extensive four day training procedure as was used in Test II. Second, the intent of virtual fixturing is to provide operators with a generic toolbox of perceptual overlays which can be implemented when confronted with a challenging task upon first encounter. Thus, taking performance measurements after extensive training on a given fixture would not reflect the intended application. Hence subjects were tested with only minimal exposure to each of the fixture designs.

The final session of Test III was a baseline test in which subjects performed the task once again without the aid of virtual fixtures. These trials were performed to allow for a comparison between this new data and the data collected in the final session of Test II. This comparison was performed in order to quantify how much (if any) performance improvement occurred during Test III as a result of learning the base task. As it turned out, subjects showed no significant performance improvement in this baseline session when compared with data collected in Test II. This result confirmed that subjects were fully trained

on the basic task (without virtual fixtures) by the end of Test II and that learning effects did not contribute significantly to the performance improvements measured during Test III.

5.4.3 Post Testing Interviews

After the completion of each of the five virtual fixture sessions of Test III, subjects were given post testing interviews in which they were asked to discuss their experience during the trials. The intent of these interviews was to gain insight into why the certain fixture designs may have been successful in enhancing operator performance and why others may not have been. Subjects were not allowed to see quantitative results of their performance before the interview session to avoid any influence on their impressions of task performance.

5.4.4 The Virtual Fixture Models:

Haptic virtual fixtures were generated and displayed to subjects using the FREFLEX exoskeleton system to impart forces. Fixture models were computed by the MERLIN VME processor shown previously in Figure 5.4. The fixture models were represented as simple functions which take as input the Cartesian position of the FREFLEX hand grip in three dimensional space. The fixture models return as output a three-dimensional force vector fed to the FREFLEX VME controller. It should be noted that the virtual fixture functions *only* depend on parameters generated in the operator space, and are therefor completely independent of feedback from the remote worksite. Thus, virtual fixture models use the location of the operator's hand to generate the resulting percept, regardless of whether or not the slave system has actually achieved the commanded position. This use of *command parameters* rather than *execution parameters* guarantees that virtual fixtures will not suffer from temporal distortions resulting from communication delays or full system bandwidth limitations. Thus virtual fixtures are generally the highest fidelity haptic percepts in the operator's perceptual reality.

Because there were many virtual fixtures models used in the testing session and because these models were varied from trial to trial, the MERLIN VME processor also received input from the TASK CONTROLLER computer which specified which fixture model should be displayed during each trial. This handshaking link between computers allowed for the rapid and automatic switching of fixture models. Because abrupt force transitions could occur when a fixture model was suddenly switched, a smooth ramping function was included to gradually fade-out one fixture model and fade-in the next. This avoided the display of disconcerting haptic jolts between trials.

Five types of virtual fixture models were implemented in this empirical study. These fixtures are referred to as: *Intersecting Surfaces, Uniform Damping Fields, Variable Damping Fields, Snap-To-Planes,*

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

and Snap-To-Lines. These particular fixture elements were *not* chosen to necessarily be a representative set of the most effective virtual fixture designs for the given peg insertion task. Rather, the particular set of virtual fixtures used in this study was chosen to simply reflect a variety of distinct and unique fixture elements. The following sections describe each of these virtual fixture designs as well as the computational models which produced them.

5.4.4.1 Intersecting Surfaces

Although these virtual fixture elements are called *surfaces*, they are very different from surfaces in the real world because they do not impose rigid constraints upon the physical domain. In fact, the mathematical model which represents these virtual surface fixtures more closely describe soft compressive springs than they do rigid surfaces. Hardware limitations made it impossible to generate stiffness values within orders of magnitude of that required by an accurate physical model of a rigid surface. Thus a perceptual approximation of a rigid surface was used rather than a physical representation. As described in Chapter 4, an effective perceptual substitute for a rigid surface percept can be produced by using spring-damper model with high viscous damping and moderate linear stiffness. Although virtual fixtures generated using such a model cannot provide the same rigid constraint upon the operator's hand motion as would a physical rigid surface, these perceptual models do present subjects with a convincing perceptual illusion of contacting a rigid surface. The particular mathematical model used in this study to haptically represent of a rigid surface is depicted below in Figure 5.7

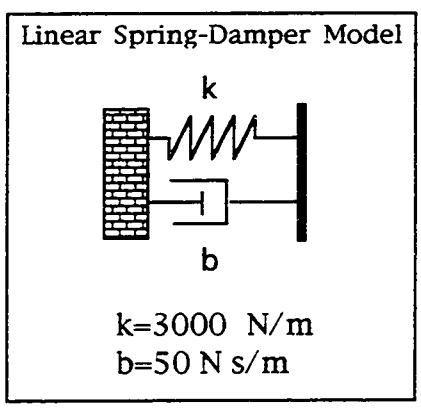

Figure 5.7: Virtual Surface Model

Having described the mathematical model for generating an individual virtual surface, the next step is to describe how multiple surfaces were combined to compose the complete virtual fixture designs. As shown in Figure 5.8, these virtual fixtures were generated by combining planar surfaces at various locations and orientations with respect to the task board. Because three different hand motions (horizontal, diagonal, and

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

vertical) were required of the subjects, three different fixture designs were used for the these trials. Figure 5.8 shows an isometric-view and a front-view of each of these fixture designs as well as an arrow which demonstrates the direction of peg motion from the start-hole from the target-hole.

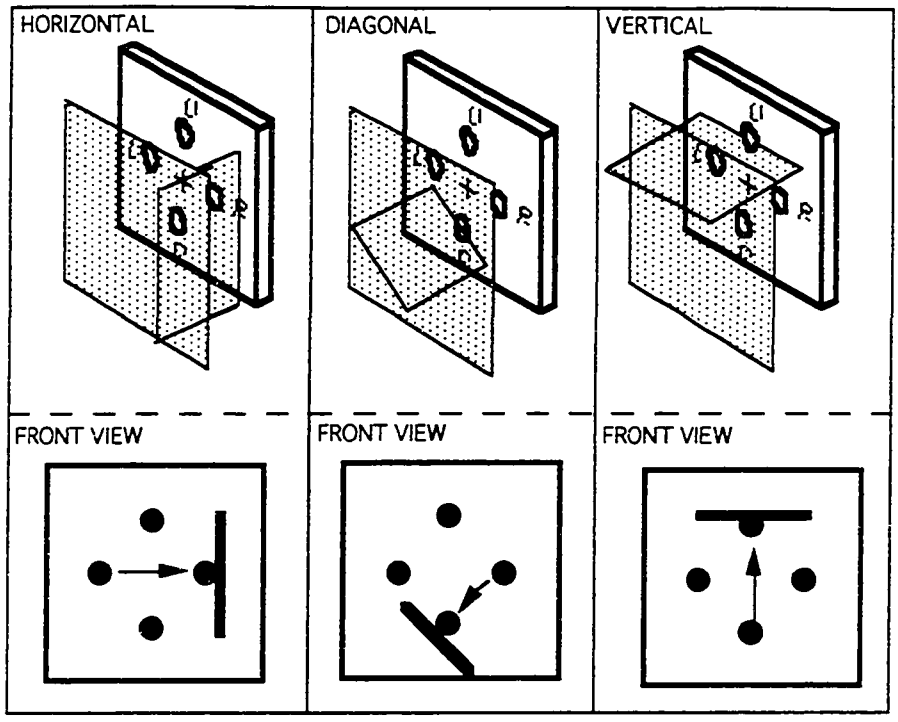

Figure 5.8: Intersecting Surface Virtual Fixtures

As shown in Figure 5.8 above, each of the three *intersecting surface* fixture designs (one for each peg motion) is drawn as overlaid on top of the task board. Although these fixtures are shown as graphical overlays above, it should be noted that they are only perceived as haptic overlays by the subjects. Each of these virtual fixtures is composed of two intersecting haptic surfaces such that one surface, to be referred to as the *back-plane*, is oriented parallel to the task board at a 3 inche distance from the board face. The other virtual surface, to be referred to as the *target-stop*, is oriented orthogonal to the first plane and located directly behind the target hole. When performing the task, the operator starts with the peg in the start-hole. Thus, the operator's hand starts at a location between the task-board and the back-plane virtual surface. When executing the insertion task, the operator will pull the peg back from the board and move the peg towards the target-hole. If the peg is pulled back too far from the task-board, the operator will encounter the haptic sensation associated with the back-plane virtual surface. If the operator overshoots the target hole, he or she will encounter the haptic sensation associated with the target-stop virtual surface. Thus these virtual

surfaces are not *guides* to the operator's hand motion, but do act as haptic cues to indicate overshoots and inefficient trajectories.

5.4.4.2 Uniform Damping Fields

This fixture is composed of a uniform viscous damping field that impedes motion in three dimensional space with an opposing force proportional to the Cartesian velocity of the operator's hand. Unlike the intersecting surface fixtures which required three different designs as described above, the uniform damping field fixture required only one fixture model for all three task motions. The particular model used in these experiments to generate the viscous damping is shown below in Figure 5.9.

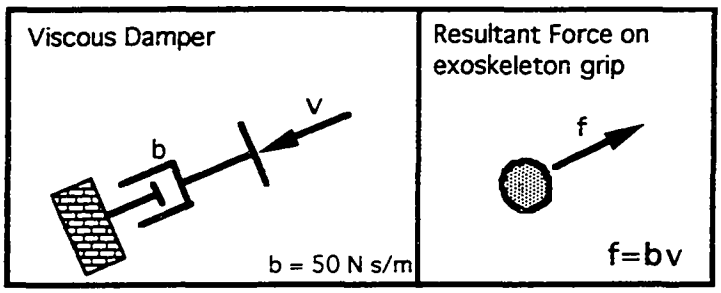

Figure 5.9: Fixture Model for Viscous Damping

5.4.4.3 Variable Damping Fields

This fixture design is similar to the Uniform Damping Field described above in that viscous damping is used to impede the rapid motions of the operator's hand. The only difference is that the Variable Damping Field fixture has a damping coefficient that varies with the spatial location of the operator's hand. For each peg motion, the variable damping field was configured so that when in the start hole, the damping coefficient was very small and impedance to the operator's motion was negligible. As the operator approached the target hole, the damping coefficient increased linearly with "closeness" (1/distance) to the target hole. When the operator's hand got to within 5 millimeters of the target-hole, the damping coefficient saturated at a maximum value of b = 75 N s/m. The resulting haptic percept was such that when the operator's hand left the start hole, there no perceptible impedance to hand motion. As the operator approached the target hole, a very strong viscous damping developed. Subjects often described the sensation as if their hand was sinking deeper and deeper into a thick liquid as they approached the target hole.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

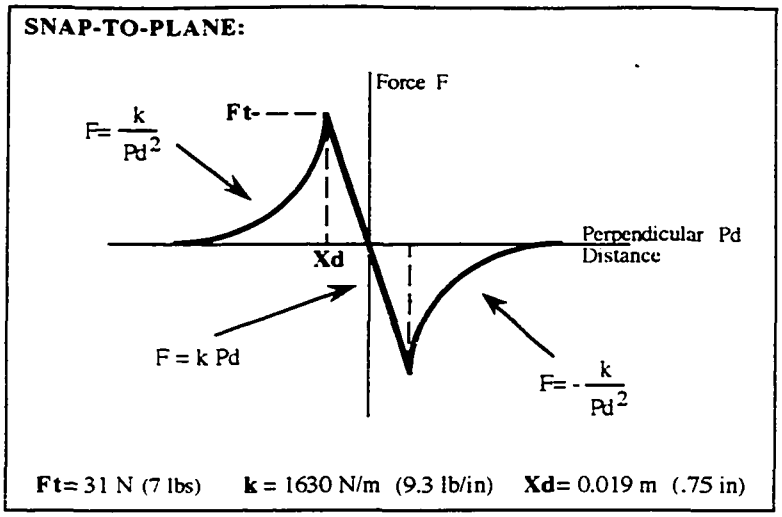

Figure 5.10: Snap-To-Plane Force Profile

5.4.4.4 Snap-To-Planes

Unlike the virtual fixture models thus far described, the snap-to-plane virtual fixture does not have a real-world analog with which to describe the sensation. In essence, the snap-to-plane is the inverse of a haptic surface. While a haptic surface is mathematical model in which an *repulsive* force is generated as one penetrates a planar region, a snap-to-plane is a model in which an *attractive* force is generated as one moves away from a planar region. While moving along the surface of the snap-to-plane, no force is applied to the operator's hand. As one deviates from the planar surface, an attractive force is applied which is linearly proportional to the deviation and thus coerces the operator's hand back into the plane. Thus the effect is like having the operator's hand attached to the planar region by a linear spring.

Of course an attractive plane would not make for a useful virtual fixture element unless there was a smooth method of entering and exiting the attractive grasp of the plane. In other words, there needed to be a means of achieving the "snap-to" aspect of this virtual percept. Perceptually, this aspect of the fixture was designed to feel like a *detent* such that the operator could *snap-in* and *snap-out* of the attractive grasp of this plane by overcoming a threshold force. Through informal subject testing and progressive design iterations, a model which provided such a force profile was developed through the combination of a central linear stiffness relation with an inverse square function shown in Figure 5.10. The graph shows force (F) wersus perpendicular distance (Pd) from the plane such that a positive force is attraction to the plane from one direction, and a negative force is attraction to the plane from the other direction. It should be noted, that all forces are oriented perpendicular to the plane and thus do not impede operator motion parallel to the plane. As a user's hand approaches the plane, an attractive force quickly ramps up with an inverse square

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

relation. Once the user comes within perpendicular distance Xd of the plane, the force profile stops increasing and sharply drops off linearly as the plane is approached. The perceptual result is a brief instant of being sucked towards the plane which smoothly dissipates to zero force when the plane is reached. This profile creates the sensation of the snap-to. Once inside the virtual fixture, if the user pulls away from the plane a linear restoring force is produced. If the user pulls away from the plane with enough force to overcome the threshold **Ft**, the force profile rapidly drops off and the operator will snap-out of the plane. The sensation has been described by subjects as "popping-in" and "popping-out" of a smooth planar channel. When inside the planar channel, the user gluides freely in two dimensions.

Having described the mathematical model used to generate a snap-to-plane, the next step is to describe how this haptic percept was used to compose complete virtual fixture designs. For each of the three peg insertion motions used in these experiments (horizontal, diagonal, and vertical) a snap-to-plane was generated along a plane coming out of the task board and passing through the start-hole and the target-hole. Thus either a horizontal, diagonal, or vertical snap-to-plane was implemented depending upon which motion was required of the subject for the given trial. In addition to these snap-to-planes, linear spring elements were added behind each target-hole. These springs were encountered when the operator's hand overshot the target-hole. Figure 5.11 shows a schematic representation of these virtual fixture designs.

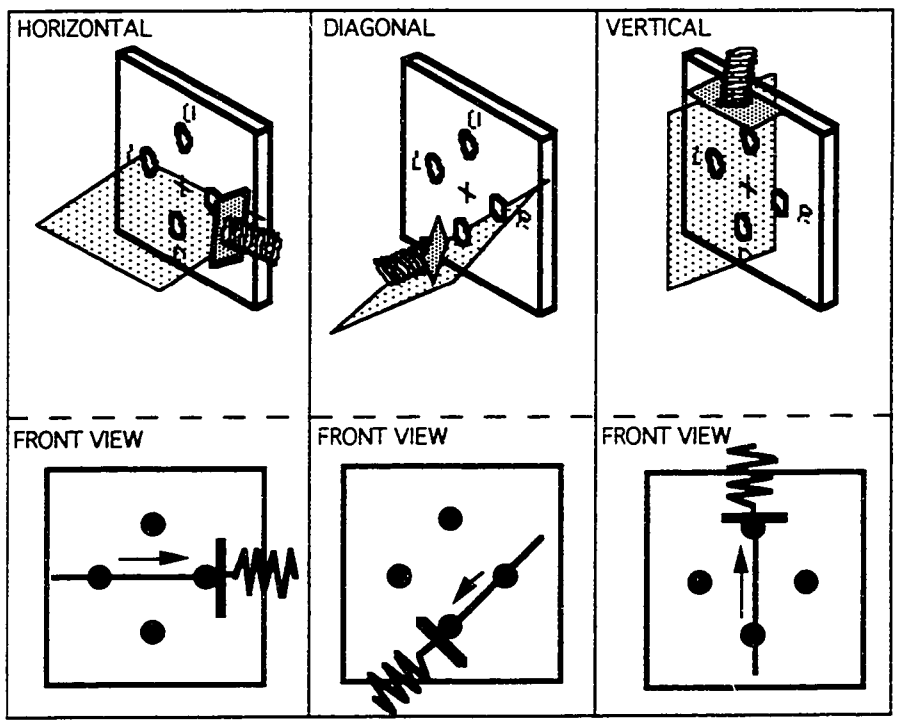

Figure 5.11: Snap-To-Plane Virtual Fixture Designs

5.4.4.5 Snap-To-Lines

The snap-to-line is essentially what the name implies, it is a line to which an operator's hand can "snap" to. The user's hand is attracted to the line radially and free to move along it axially. Snap-to-lines are very similar in principle to the snap-to-planes described above. In fact, the mathematical model which describes a snap-to-line is identical to that shown in Figure 5.10, with the one distinction: the radial distance from a line is used as the independent variable instead of the perpendicular distance from a plane. When an operator's hand gets within range of the snap-to-line, it is attracted to the line by a quickly ramping radial force. Once the snap-to occurs, operator is free to move axially along the line, but it is restricted from deviating off the line by a linear force profile. It is as if a linear tension spring connects the operator's hand to the line. The operator can break free from the attractive influence of the snap-to-line by exceeding the threshold force. For these tests, the threshold force was set to 7 lbs.

Having described how a snap-to-line percept is generated, the next step is to describe how they were implemented as virtual fixtures. As shown in Figure 5.12, the snap-to-lines were used in very simple virtual fixtures each of which was composed of a single snap-to-line passing through the target-hole and oriented perpendicular to the task board. These fixtures were intended to aid in target acquisition by snapping the operator's hand in line with the target-hole.

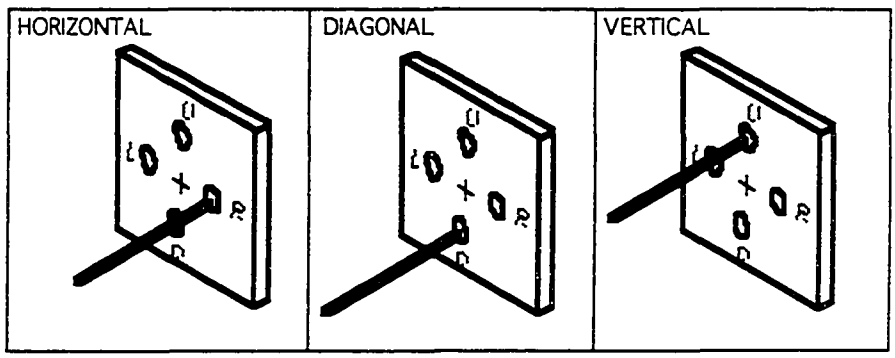

Figure 5.12: Virtual Fixture Designs Using Snap-To-Lines

5.5 RESULTS

Movement times were recorded for 30 peg-insertion trials during each of the seven data collection session performed by all subjects. These 30 trial sessions were composed of three groups of 10 trials. Each of these groups of 10 trials represented the movement time results for one of the three peg motion directions (horizontal, diagonal, and vertical). Means and standard deviations were computed for each of these peg motions. Thus the results of each session consisted of three mean movement times and three standard deviations for each of the seven subjects. This data was further distilled by computing the mean movement times and standard deviations across all 7 subjects. These mean values are shown below in Figure 5.13.

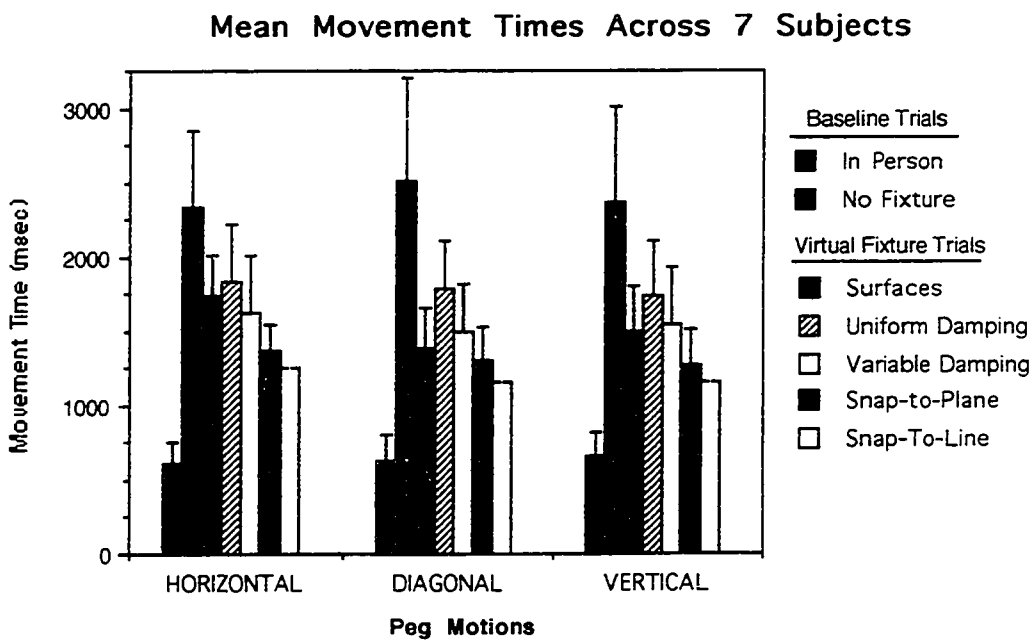

Figure 5.13: Movement Time Results Across all Test Sessions

The above figure is divided into three groups of data, each one representing one of the three peg motions performed by the subjects. Each of these groups of data is further broken down into columns, each of which represents the summarized results from a single testing session. The first column of each group represents the initial baseline session in which subjects performed the task in-person while standing in front of the task board. The second column of each group represents the baseline session in which subjects performed the task remotely while using the exoskeleton system but not using any virtual fixtures.

The final five columns of each group represent the trials performed using each of the five virtual fixture designs presented during this study.

This data was then further compiled by computing means across the three peg motions performed by each subject. These total means are shown below in Figure 5.14. The values reflect the overall performance across three different peg motions.

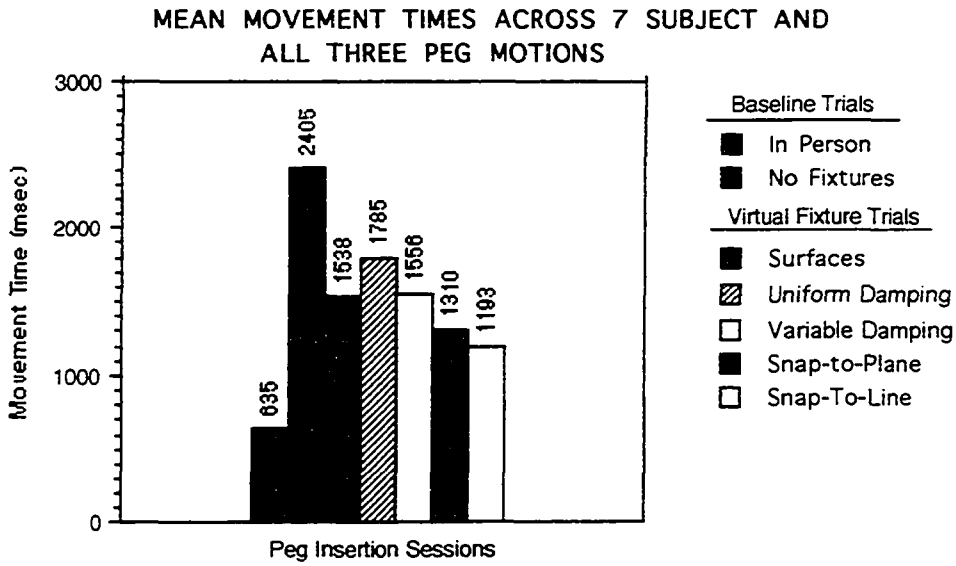

Figure 5.14: Mean Movement Time Results Across all Test Sessions

5.6 ANALYSIS

When comparing human performances in Fitts Law style peg-insertion tasks as in Chapters 2 and 3, one usually derives an Index of Performance (I.P.) which is computed as the difficulty of the task (in terms of Index of Difficulty) divided by the recorded movement time. Index of Performance is a valuable parameter for standard psychophysical testing because it can be compared across experiments as a normalized performance measure. Nevertheless, I.P. is not that useful for human tests which do not require comparison outside of a single set of experimental trials. What is more, I.P. is not a particularly insightful measure when computing human performance in a telepresence situation because it does not reflect the performance degradation due to the telepresence link. Thus, for this set of experiments, a new measure of human

performance was introduced to be referred to as the Index of Telepresence Performance (I.T.P) which is computed as follows:

I.T.P. = Movement Time in Person / Movement Time in Telepresence Task

By computing such a measure, we get an index which reflects telepresence performance as a percentage of a natural, in-person baseline. An I.T.P. of 1.0 means that performance was equally good in the telepresence system as in person. An I.T.P of less than 1.0, reflects a performance degradation and an I.T.P of more than 1.0 reflects a performance enhancement above natural ability. For this set of tests, I.T.P was computed for all test session and are shown below in Table 5.3. Since these performance values are normalized by the performance in the in-person task, they can be compared across the three task motions (horizontal, diagonal, and vertical) without differences in task difficulty effecting the results.

Table 5.3: Index of Telepresence Performance (I.T.P.) for all Sessions:

	In Person	Telepresence No fixture	Two Surfaces	Uniform Damping	Variable Damping	Snap-To Plane	Snap-To Line
Horizontal	1.00	0.26	0.35	0.33	0.37	0.45	0.49
Diagonal	1.00	0.25	0.46	0.36	0.42	0.49	0.54
Vertical	1.00	0.28	0.44	0.38	0.43	0.52	0.57
MEAN	1.00	0.26	0.42	0.36	0.41	0.49	0.53

Although the values in the above table demonstrate that performance degradation in the no-fixture telepresence session was significantly greater than the performance degradation when virtual fixtures were used, a final level of data processing was required to provide a direct comparison of each virtual fixture session to the no-fixture telepresence baseline. In order to process the mean values of Index of Telepresence Performance into a form which would directly represent the effect that virtual fixtures had upon operator performance, the *percentage performance increase* with respect to no-fixture teleoperation was computed for each fixture design. These values, representing the percent increase in I.T.P for each virtual fixture trial with respect to the no-fixture baseline, are shown below in Table 5.4.

Table 5.4: Percent Increase In I.T.P. For Each Virtual Fixture Session

	Two Surfaces	Uniform Damping	Variable Damping	Snap-To Plane	Snap-To Line
Horizontal	34 %	27 %	44 %	72 %	87 %
Diagonal	82 %	41 %	68 %	93 %	116 %
Vertical	58 %	36 %	52 %	87 %	103 %
MEAN	59 %	35 %	55 %	84 %	102 %

Although the above results clearly demonstrate the effect that virtual fixtures have upon raw operator performance, they say little about the effect that fixtures had upon operator consistency. To address this issue of consistency, the coefficient of variation (COV) was computed for all sessions. The COV is an indicator of operator consistency which reflects the variation across the trials of each session as a percentage of the mean value for that session. The mean COV for each session is shown below in Table 5.5. It should be noted that a number training periods preceded each of the two baseline sessions but no training was given prior to the fixture sessions. Thus, it would be expected that fixture performance be more variable than baseline performance. As is shown in the table, performance in all of the fixture sessions were actually more consistent than the baseline sessions. These results suggest that operators can implement new virtual fixtures when confronted with a given task and not have performance consistency suffer even though they were not specifically trained on that fixture design. This lends support for the notion of an interactive perceptual workstation in which operators design and implement their own fixture designs when confronted with a task.

Table 5.5: Mean Coefficient of Variation for Each Session

	In Person	No fixture (remote)	Two Surfaces	Uniform Damping	Variable Damping	Snap-To Plane	Snap-To Line
Horizontal	0.24	0.23	0.16	0.19	0.24	0.13	0.16
Diagonal	0.28	0.27	0.21	0.24	0.21	0.18	0.20
Vertical	0.24	0.27	0.20	0.22	0.25	0.20	0.16
MEAN	0.25	0.26	0.19	0.22	0.23	0.17	0.17

5.7 DISCUSSION

5.7.1 In Person Trials Compared to Baseline Telepresence Trials

Before addressing the results of the five virtual fixture test sessions, the results of the two baseline sessions should be compared to provide insight into the effect that the telepresence link has upon raw operator performance. As shown in Figure 5.14, the mean movement time in-person trials was 635 msec while the mean movement time for trials performed remotely (though the use of the telepresence interface) was 2405 msec. This is an increase in movement time of almost 300%. If we look at this in terms of Index of Telepresence Performance as shown in Table 5.3, we can say that average operator performance using the telepresence interface was 26% of natural manual performance. These raise the following question: Why is operator performance so greatly degraded as result of task execution through the telepresence system?

Before addressing this question, the reader should be reminded that the in-person trials were performed through the use of the *exoskeleton encumbrance simulator* which emulated the physical encumbrance of the telepresence link during the baseline in-person trials. Thus, the physical encumbrance imposed by telepresence system should not have contributed significantly to the performance degradation. The reader should also be reminded that the stereo-vision system used in this experiment was composed of direct binocular optics focused on the remote site. This vision system provided very high fidelity visual feedback and is therefor not likely to have contributed to the performance degradation in the telepresence trials. Finally, it should be pointed out that the maximum time-delay in issuing commands to the remote manipulator from the exoskeleton was on the order of 6 ms. This lag is considered inperceivable to a human operator (in general, people can not perceive time delays below about 60 ms). Hence, time-delay

was probably not the source of the performance degradation Limited bandwidth of the remote manipulator is one possible source of performance degradation, but the physical encumbrance of the exoskeleton seemed to act as a fairly effective impedance match between the operator and the remote manipulator system. As a result, no subjects reported feeling as though the slave robot was having trouble keeping up with their input commands. Thus it is unlikely that manipulator bandwidth was the primary source of performance degradation.

So why the huge performance degradation in the telepresence trials when compared with the in-person baseline? Although some of the drop in performance could possibly be attributed to physical hardware limitations such as the bandwidth, range of motion, and encumbrance of the master-slave system, every attempt was made to minimize the effect of such parameters. Other than these unavoidable limitations in the man/machine interface, the only other significant difference between the in-person and the telepresence trials was the basic fact that the in-person trials were performed with *direct human presence* in the worksite while the telepresence trials were performed with *simulated presence* projected into the worksite from a remote location. Thus a likely explanation for the drop in performance is that the simulated presence provided by the telepresence link was inferior to the natural presence provided by direct human contact. Sheridan [1992] suggests that task performance as well as training efficiency are parameters which are highly dependent upon an operator's sense of presence in a remote or virtual environment. Thus it does seem reasonable that reduced operator presence could be a primary source of performance degradation in the telepresence trials.

Although reduced operator presence is a reasonable explanation for the recorded drop in subject performance, how can we account for such a reduction in operator presence? Held [1990] suggests that the conditions which contribute to operator presence in a remote or virtual environment include *the sensor's ability to discriminate figure from background*; *the perception of the sensor's independence from the remote or virtual environment* (i.e.. sensory content varies with positional changes of the sensor); and *the ability of the sensor to access the world at will* (i.e.. the ability to move the eyes or limbs to explore environment from different perspectives). Sheridan [1992] refined these criteria by proposing that the determinates of presence are: the *extent of sensory information* (transmitted bits of salient sensory feedback), the *operator's control of the relation between sensors and the environment*, and *the operators ability to modify their physical surroundings*. Loomis [1993] proposed that presence is dependent upon a "lawful relationship" between an operator's efferent commands (i.e.. kinesthetic motions), and the afferent sensory information reflected from the environment. He went on to say if such a relation between afference and efference does not exist or is too complex for the operator's nervous system to represent, the illusion of presence will be degraded.

In light of these recent theories about the contributing factors to remote and virtual presence, operator presence may have been reduced in this set of experiments as a result of limitations in the *extent of sensory information* reflected to operators, limitations in the operator's control over *the relation of sensors to environment*, and/or limitations in the operator's *ability to modify the physical environment.*. Considering the last case first, it is unlikely that limitations in the operators ability to modify the physical environment had any effect upon operator presence in this set of experiments. Although the remote environment consisted of a fairly sparse worksite including just a task board and pegs, operators maintained complete control over the remote manipulator and could move pegs at will. As for limitations in the extent of sensory information reflected from the remote environment, it is possible that such limitations in this experimental set-up had some effect upon operator presence. Although visual information reflected from the remote environment consisted of high fidelity stereo images, the field of view was reduced as compared to natural human vision. While the audio information reflected from the remote environment consisted of high fidelity sound, the audio feedback was devoid of binaural content. And while much effort went into providing haptic feedback from the remote environment, the fidelity of such haptic cues were limited by the bandwidth of the exoskeleton system.

Although limitations in the sensory feedback from the remote environment may have contributed somewhat to degradation in operator presence, the experimental setup was designed to provide a rich sensory link to the remote site. Thus it seems unlikely that such dramatic performance degradations could be attributed to the limitations in sensory feedback mentioned above. It seems more likely that the dominant factor in reduced presence was the limited modalities through which afference could be related to efference in the telepresence link. Whether we present it in terms of the requirement for *a lawful relation between afference and efference* as proposed by Loomis, or if we present it in terms of the necessary *sensor independence from the remote environment* as maintained by Held and Sheridan, the reduced sense of presence in the baseline telepresence trials was likely the result of the limited modalities through which the operators could modify the incident sensory feedback from the remote environment.

Because this experiment consisted of small amplitude peg-insertion trials, head rotations or body shifts were not required to comfortably execute the task. Thus both head and body were held fixed during telepresence testing. Such simplifications are very common in telepresence environments to reduce the hardware requirements for a given task. Because this experimental task was performed in a small remote workspace which fit well within the operator's field of view, there seemed to be little justification for implementing expensive and complex hardware such as a head tracking camera system and a mobile exoskeleton platform. Although such mobility was not required to enable the peg-insertion task, the mobility may have been essential for achieving a strong sense of presence within the remote site. This experiment only implemented hardware which seemed necessary and relevant for the execution of the task at

hand. The primary modality through which afference could be tied to efference was the mapping between arm motions and visual feedback of robot position. There were also brief instants at the beginning and end of task execution when the peg came in contact with the board. At such instants there was also a matching between kinetic motions and haptic feedback. In any case, because head motion and body shifts were restricted, the modalities through which afference and efference could be related were sparse. It seems highly possible that such activities, whether they are relevant to actual task execution or not, may be involved in the construction of a rich mental model of the remote environment and thus important to operator localization with respect to the task site. Many researchers have demonstrated that externalization and localization of 3D audio images is highly dependent upon head motions even though pinna related transfer functions can theoretically resolve such audio cues without head motions (e.g., Young, 1931; Wallach, 1940; Durlach 1992). Although such audio effects were not relevant to this study, they may lend credence to the general hypothesis that there is more to be considered when designing a telepresence man/machine interface than is directly implied by the requirements of a given task. This is yet another facet of the Design-for-Perception methodology.

5.7.2 Virtual Fixture Trials Compared to No-Fixture Baseline Trials

Turning next to the virtual fixture peg-insertion trials, Figure 5.15 clearly demonstrates that the movement times recorded in all virtual fixture sessions were greatly improved over the no-fixture baseline. Table 5.4, which lists the percent increase in the Index of Telepresence Performance (I.T.P) for each of the virtual fixture sessions as compared to the no-fixture baseline, demonstrates that all virtual fixtures tested resulted in significant performance improvements over the no-fixture trials ranging from 35% to 102% in the five virtual fixture sessions. The mean performance improvement across all virtual fixture trials was 67%. Clearly, virtual fixtures were highly effective in enhancing operator performance in the given peg-insertion telepresence task.

Considering that the most reasonable explanation for the recorded degradation in telepresence performance (as compared to in-person performance) was reduced operator presence at the worksite, it seems reasonable to attribute at least part of the performance improvement resulting from virtual fixture use to an increase in operator presence. Considering that the most likely explanation for reduced operator presence in baseline trials was the reduced modalities through which afference could be related to efference, the role of virtual fixtures in enhancing the illusion of *presence* is highly supportable; Since virtual fixtures introduce perceptual information into the operator's perceptual reality, they provide additional pathways through which afference and efference can be related. Consider the most generic of all virtual fixtures tested in this set of experiments, the uniform damping field. By overlaying a simple damping field on top of the operator's perception of the remote environment, the operator's afferent kinesthetic commands are no longer

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

only related to visual feedback from the remote environment, but are also related to haptic feedback provided by the perceptual overlay. The same can be said of the other four virtual fixtures tested in this set of experiments. Thus, virtual fixtures may stimulate operator presence by providing additional sensory modalities through which afference and efference can be related. This in turn, may contribute significantly to the performance increase associated with fixture use.

Beyond increasing the number of modalities through which afference could be related to efference, virtual fixtures are likely to enhance performance through other means. The following paragraphs address each of the five virtual fixtures independently and proposes additional explanations for the performance improvement resulting from their use. In addition to the quantitative test results, post-testing interviews and recorded hand trajectories are discussed to provide deeper insights into the effects of virtual fixtures.

5.7.2.1 Intersecting Surfaces

Intersecting surface fixtures consists of two planar haptic surfaces positioned with respect to the task board as shown in Figure 5.8. These fixtures produced a 58% mean performance improvement as compared to the no-fixture trials. Although referred to and perceived as "surfaces", the mathematical model which generated these haptic percepts were far from the rigid kinematic constraints which the name "surface" implies. As shown in Figure 5.7, these virtual surfaces were modeled more like compressive springs than rigid walls. Thus these fixtures did not act as rigid kinematic constraints upon the operator's kinesthetic activities. Rather, these virtual fixture elements functioned as haptic overlays with more of a perceptual than physical impact upon task execution.

Each of these fixture designs included one surface oriented parallel to the task board (the back-plane) and a second surface oriented orthogonal to the task board (the target-stop). When performing the task, the operator first encountered the *back-plane* when pulling the peg out of the start-hole. The operator would contact this surface when the peg had been fully removed from the hole and the tip was well clear of the task-board. During post testing interviews, most subjects reported using this haptic cue to supplement or replace the visual feedback indicating that the peg had been pulled far enough out of the task-board that it could be moved safely towards the target-hole. This motion component of task execution was a major concern for subjects, since if the peg was not pulled back far enough, the peg would get stuck in the hole and waste execution time. If the peg was pulled back too far, arm trajectories were unnecessarily large and inefficient, again wasting time.

A similar effect occurred during the insertion phase of the telemanipulation task. The *target-stop* haptic surface was oriented orthogonal to the task-board and was encountered when the operator positioned

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

the peg in front of the target-hole. During post testing interviews, most subjects reported using this haptic cue to supplement or replace the visual feedback indicating that the peg was in position and ready for insertion. This component of task execution was a major concern to subjects because operators very often overshot the target-hole when trying to approach it as fast as they could. It is likely that haptic cues provided by the back-plane and target-stop surfaces increased the operator's processing capacity by providing an additional sensory modality through which successful peg removals and insertions could be assessed and reacted to. In addition, since human reaction time to haptic stimuli is faster than reaction time to visual stimuli , the incorporation of the haptic sensory information into task execution not only added a useful sensory modality, it added a control pathway with higher bandwidth.

5.7.2.2 Uniform Damping

Although this virtual fixture design also introduced haptic information into the user's perceptual reality, the resulting percepts were very different than those produced by the virtual surface models. Rather than providing haptic cues based on physical location of the operators hand, uniform damping fixtures generated haptic information based on the *dynamics* of hand motion. Thus, the uniform damping virtual fixture design did not provide discrete haptic cues as did the surface fixtures. Rather, the damping fixture provided an environment rich in continuous haptic feedback. As shown in Table 5.4, the mean performance improvement resulting from use of uniform damping virtual fixtures was 35% as compared to the no-fixture trials. This 35% performance improvement is very interesting result in light of the fact that this fixture was a generic design, having no parameters which reflect the given task or the given remote worksite. Such a uniform damping virtual fixture could be overlaid at the touch of a button on top of *any* worksite, with no preplanning or calibration. The effectiveness of such a generic perceptual overlay is a strong confirmation that an interactive perceptual workstation could be developed which allows fixtures to be rapidly designed from generic building blocks.

The question still remains, why did the uniform damping fixture design enhance manual performance in the given peg insertion task. One possible reason was simply that the additional of haptic feedback provided an alternate sensory pathway through which task execution could be monitored. Proper execution of the given peg-insertion task without the use of a damping fixture was primarily maintained by visual feedback (i.e. proper execution of the task looked a certain way). With the damping perceptual overlay, proper execution not only looked a certain way, it *felt* a certain way. This presumably allowed subjects to monitor hand dynamics using direct haptic feedback as a perceptual cue. In post testing interviews, all subjects attributed the performance improvement associated with the uniform damping fixture with a feeling of enhanced control of the remote manipulator. In a sense, adding artificial damping to the environment filtered out the subject's high frequency motions and allowed them to control the remote

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

manipulator with reduced jitter and control oscillations. This was particularly useful when wielding the increased inertia of the exoskeleton because such added inertia results in a slower decay of human control oscillations. For example, without the damping fixture most subjects overshot the target hole and oscillated a few times before inserting the peg. With the damping fixture, such oscillations were greatly reduced. To demonstrate this effect, sample hand trajectories were recorded from each subject. Figure 5.15 shows a typical hand trajectory performed without the aid of any virtual fixtures. As you can see, there is a characteristic oscillation present during the target acquisition phase of the insertion task. This oscillation can be attributed to both the increased inertia of the exoskeleton system as well as the reduced bandwidth of the control pathway. Figure 5.16 shows a trajectory performed by the same subject when the uniform damping field is overlaid on top of the remote worksite. In this case, the oscillation is greatly reduced.

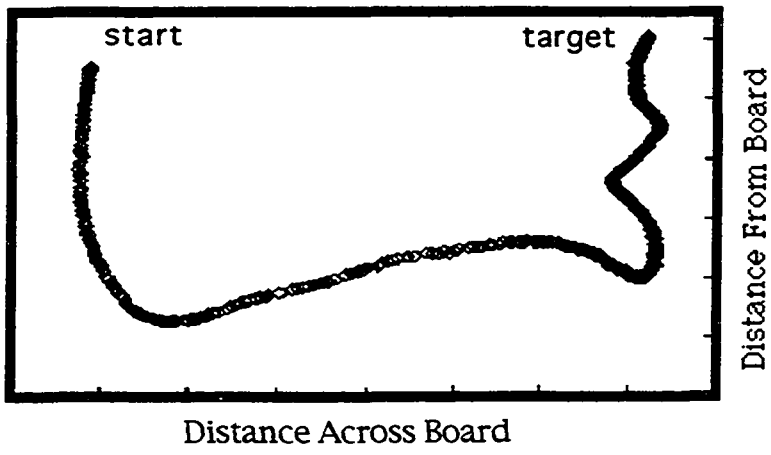

Figure 5.15: Peg Insertion Trajectory with No Virtual Fixture

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

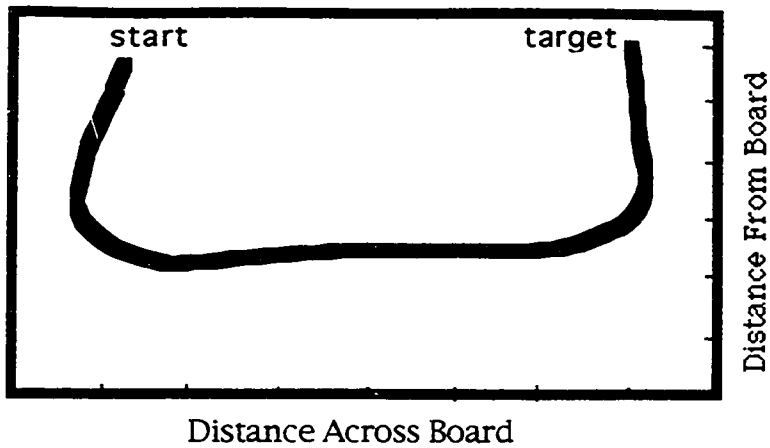

Figure 5.16: Peg Insertion Trajectory with Uniform Damping Field

5.7.2.3 Variable Damping

Because the uniform damping fixture seemed to contribute most to operator performance during the final target acquisition (insertion) stage of the peg-insertion task, a *variable damping* fixture was designed which increased the coefficient of damping as the operator's hand approached the target-hole. When the operator pulled the peg out of the start-hole and began a ballistic approach towards the target-hole, the damping was very small and therefore induced very little impedance upon hand motion. As the operator approached the target, the virtual impedance ramped-up linearly and provided maximal damping during the target acquisition phase of the insertion task. Thus maximum damping was applied at the point in the task when overshoot was most detrimental to performance. As listed in Table 5.4, the mean performance improvement resulting from use of the variable damping fixture (as compared to the no-fixture case) was 55%. Not only did this fixture provide a significant performance improvement over the no-fixture case, it provided a substantial improvement over the uniform damping fixture design. This result further substantiates the hypothesis that the virtual damping enhanced the operator's manual control during the target acquisition phase of the insertion task. Post testing interviews revealed that all subjects preferred the variable damping fixture to the uniform damping fixture.

5.7.2.4 Snap-To-Plane

As shown in Figure 5.10, the snap-to-plane virtual fixture is essentially a smooth planar region which an operator can "snap into" and "snap out of" by overcoming a small threshold resistance. This is a particularly interesting fixture to test on human subjects because it is a haptic percept which has no real-

world analog. Unlike a haptic surface virtual fixture, a snap-to-plane is such that motion along the defined plane corresponds to a zero resultant force, and deviations away from the plane corresponds to a positive restoring force. Thus a snap-to-plane can almost be looked at as the inverse of a haptic surface model As shown in Figure 5.11, snap-to-planes were used in combination with virtual springs to compose a compound virtual fixture for the given task.

Despite the fact that snap-to-planes are abstract haptic constructs which posed a new perceptual experience for the test subjects, results revealed a dramatic performance improvement as a result of such fixture use. As shown in Table 5.4, an 84% mean performance improvement resulted from use of the snap-to-plane virtual fixtures as compared to the no fixture case. A likely reason why the snap-to-plane virtual fixture so greatly enhanced operator performance is that it maintained the user's trajectory in the plane of the task and thus reduced the dimension of the task space by one degree of freedom. Thus the snap-to-planes acted to simplify the task and reduce the mental processing required of the subjects. It should be noted that operators could easily snap-in or snap-out of the plane at will by overcoming a small threshold force. Thus although the subject could use the haptic information to simplify the task space, such virtual perceptual overlays were not actually kinematic constraint upon operator motion. All subjects reported preferring the snap-to-plane perceptual overlay to that of a virtual surface because the neutral position was a zero force resultant.

5.7.2.5 Snap-To-Line

As shown in Figure 5.12, a snap-to-line virtual fixture is essentially a line emanating from the target hole (orthogonal to the task board) to which the operator's hand can "snap". Once such a snap-to occurs, the operator can slide freely along the axis of the line but is restricted from deviating off the line by an attractive force. A mild, but distinct, jolt provides a powerful haptic cue that indicates a snap-to has occurred. Like the snap-to-plane, the snap-to-line can be exited by overcoming a small threshold force. Like a snap-to-line, a snap-to-plane is a particularly interesting virtual fixture for human testing because it employs an abstract haptic percept which has no real-world analog.

Despite the fact that snap-to-lines are abstract haptic constructs which pose a new and unusual perceptual experience for subjects, testing has revealed that the use of such virtual fixture resulted in the greatest performance improvement of all fixtures tested. As shown in Table 5.4, a 102% mean performance improvement resulted from use of the snap-to-line virtual fixtures as compared to the no fixture case. Looking at the movement time results shown in Figure 5.14, the mean time required for a peg insertion to be completed without the aid of any virtual fixtures was 2405ms. With the aid of the snap-to-

line, mean movement times dropped by more than half to 1193ms. A likely reason why the snap-to-line virtual fixture so greatly enhanced operator performance is that such perceptual overlays provided a strong haptic cue to aid in target acquisition. Without the use of such virtual fixtures, alignment of the peg with the target hole was performed as a primarily visual task that required hand-eye coordination executed through visual/kinesthetic coupling. With the snap-to-line, the task was much less dependent upon the visual modality and could be performed primarily using the haptic cue. For example, upon encountering the snap-to-line virtual fixture, a distinct haptic jolt would alert the operator that a snap-to had occurred and thus indicate that the peg was aligned with the hole and could be inserted. What is more, the snap-to-line acts as an active phenomenon in the environment which actually encourages proper peg alignment through a mild resistive force which opposes deviations from the peg alignment line.

5.8 CONCLUSIONS

As described in the discussion above, this set of subject testing has confirmed that computer generated haptic sensations can be used as perceptual overlays which greatly enhance operator performance in telepresence tasks. Despite hardware constraints such as mechanical bandwidth and computational power which limited the form, content, and fidelity of the haptic percepts presented to subjects during these experiments, virtual fixtures were constructed which enhanced operator performance by as much as 102% in the given peg-insertion tasks. The mean performance improvement across all virtual fixtures tested in this round of experiments was 67%. This mean performance improvement actually exceeded the 49% increase discussed in Chapter 2 which resulted from the use of *ideal* haptic sensations displayed using real mechanical fixtures rather than force reflecting hardware. The greater performance increases resulting from the use of computer generated haptic percepts compared to ideal haptic percepts suggests that the flexibility to produce abstract haptic sensations is more important to the development of effective perceptual overlays than is the ability to produce high fidelity representations of particular haptic percepts. Thus the potential of virtual fixturing is not limited by hardware requirements for high fidelity haptic display. This result acts as a general confirmation of the notion of Design for Perception because it demonstrates a real world situation in which physical realism in the form of haptic fidelity is not as important as the ability to embellish an operator's perceptual reality with abstract haptic cues.

In addition to confirming the basic premise that computer generated haptic percepts can be effective perceptual overlays, this set of experiments was the first empirical study to expand the virtual fixture concept to three dimensional space. The results clearly confirm that three dimensional haptic percepts can be used as effective perceptual overlays in a given telepresence task. Also, this set of experiments was the

first test of more general haptic percepts because previous experiments were limited to virtual fixtures composed of rigid surfaces. Thus while the previous tests in Chapters 2 and 3 were restricted to fixtures which could act as rigid kinematic constraints on operator hand motion, this new set of testing included a wide variety of haptic percepts which enhanced performance without imposing rigid kinematic constraints upon operator motion.

Beyond testing the effectiveness of perceptual overlays, this set of experiments also recorded two baseline sessions which allowed a comparison between direct human performance and un-aided telepresence performance. The results of such baseline trials have revealed a 300% increase in movement times when comparing telepresence to direct human activity. The dramatic drop in performance associated with the telepresence trials (which could not be attributed to encumbrance effects of the exoskeleton) were likely the result of reduced operator *presence* in the worksite. It has been hypothesized that reduced presence in the telepresence trials was the result of the limited modalities through which *efference* (kinesthetic commands) could be related to *afference* (sensory feedback from the remote site). This line of reasoning suggests that the performance improvements resulting from virtual fixtures may be partially because enhancement of the operator's perceptual reality with abstract overlays creates additional sensory pathways through which afference and efference can be related. Thus perceptual overlays may improve operator presence in a remote or virtual environment and subsequently enhance performance in manual tasks.

In addition to possibly enhancing operator presence in the telepresence worksite, virtual fixtures may have enhanced performance in a number of other ways. It is hypothesized that virtual fixtures may have enhanced performance in the given telepresence task in the following ways:

1: Virtual Fixtures may increase operator information processing capacity by increasing the number of sensory modalities to the task space. The addition of sensory information could enhance performance by allowing tasks to be monitored through multiple sensory pathways simultaneously. The addition of sensory information to the task space could also enhance performance by enabling the use of a more appropriate or more efficient sensory modality for a particular task. The addition of sensory information may also enhance performance by displacing the perceptual burden from other overburdened sensory modalities.

2: Virtual Fixtures may reduce the information processing requirements of the task by eliminating degrees of freedom from the task space. For example, by introducing a haptic percept into the task environment which explicitly informs the user that they have removed the peg from the start-hole and are at a safe distance back from the task-board, operators no longer have to waste processing capacity on assessing their distance from the task board during execution.

3: Virtual Fixtures provide a rich perceptual environment for task execution which may allow operators to generate more effective mental models of the given task. For example, without haptic perceptual overlays, proper execution of the peg insertion task can be described only as "looking a certain way". Once haptic percepts, such as a uniform damping field, are introduced into the users perceptual reality, proper execution can also be described as "feeling a certain way". By providing users with additional ways to conceptualize and analyze proper task execution, virtual fixture are likely to enhance manual performance.

Chapter Six

6. Binaural Perceptual Overlays: Audio Fixtures in 3D Space

Thus far, a series of human tests have been performed which demonstrate the effectiveness of haptic and audio perceptual overlays to enhance operator performance in telemanipulation tasks. Chapter 2 and Chapter 3 began the empirical exploration of virtual fixturing by testing simple haptic and auditory percepts in two-dimensional space. Chapter 5 expanded the testing of haptic perceptual overlays to more complex haptic constructs in three-dimensional space. The current chapter further expands the testing of virtual fixturing, by extending the initial study of simple two-dimensional audio overlays, as described in Chapter 2, to more complex *binaural* audio percepts in three-dimensional space.

The following sections describe a simple set of experiments which implement virtual fixtures composed of overlaid binaural audio signals. An *Audio Localization Cue Synthesizer* (ALCS), was employed to generate and display perceptual overlays constructed from three-dimensional audio localization cues. A Fitts' Law peg insertion test was performed to quantify the effect that such binaural perceptual overlays have upon operator performance in a telemanipulation task. Results confirm that operator performance can be enhanced by embellishing the user's perceptual reality with abstract audio localization cues. This is a particularly interesting result because synthetic binaural information represents a purely perceptual phenomena, having no direct analog in physical reality. While haptic cues can be directly related to the physical phenomenon they represent, binaural localization cues *only* have meaning in the observer's perceptual representation of reality. In other words, virtual binaural cues represent a synthesized *proximal stimulus* which does not have a direct physical relationship to an inferred *distal stimulus*. The only relation which can be drawn between proximal and distal stimuli is perceptual. Thus we can look upon this

testing of audio localization cues as a pure example of enhancing operator performance by *only* embellishing a user's perceptual rather than physical reality.

6.1 EXPERIMENTAL HARDWARE

The same telemanipulation peg-insertion paradigm was implemented in this round of empirical testing as was used in the 3D haptic tests discussed in Chapter 5. Thus, most of the experimental hardware is common between these two experiments. In the interest of brevity, only those components which are unique to 3D audio localization testing will be fully described here. For a complete description of the other hardware components, please refer back to Chapter 5.

6.1.1 Overall System Setup

This round of empirical testing required an extensive array of hardware and software to allow human operators to naturally control a robot in a remote telepresence worksite while at the same allowing virtual fixtures to be overlaid on top of the user's perception of that worksite. This setup included a FREFLEX force reflecting exoskeleton, an Air Force/Navy Teleoperator Task Board, a Merlin Robot Arm, A Stereo Vision System, and an *Audio Localization Cue Synthesizer* (ALCS). Complete descriptions of all hardware subsystems, except for the ALCS, can be found in Chapters 2 and 5. Also, Chapter 5 describes the computational architecture which allows these subsystems to function in unison. The ALCS, which is the only piece of equipment unique to the audio localization testing, is described below.

6.1.2 *Audio Localization Cue Synthesizer* (ALCS):

In 1930, Firestone first demonstrated *simulated auditory localization* by fixturing a dummy head with microphones at the base of each ear canal. The acoustic signals from each microphone were presented to a remote listener through headphones. This binaural setup was the first demonstration that three-dimensional auditory percepts could be displayed to a user through synthetic means. Since that time, much effort has been directed towards resolving the physical transfer functions used by the human perceptual system to perform 3D audio localization from binaural sensations. Perceptual parameters such as Interaural Time Delay (ITD), Interaural Amplitude Difference (IAD), and Head Related Transfer Function (HRTF), have been well characterized by the psychoacoustics community. With such perceptual information in hand, designers have begun to develop computer systems capable of artificially producing the appropriate proximal representation of a distal auditory signal located in three-dimensional space.

Developed at Armstrong Aerospace Medical Research Laboratory's Bio-Acoustics Branch (AAMRL/BBA), the *Audio Localization Cue Synthesizer* is an electronics module designed to artificially encode naturally occurring spatial information in the form of an audio signal that can be presented to a user through stereo headphones. The encoded audio signal can be presented such that it is perceived as originating from a particular location in three-dimensional space. As the listener changes the position and orientation of his/her head, the relative position of the virtual sound-source is updated to make the sonic percept appear stationary with respect the listener's environment. As currently implemented, the ALCS can present a single sound source such that the azimuth, elevation, and distance of the signal can be specified by computer control. Thus, the ALCS provides complete 3D control over the perceived spatial origin of a given audio signal.

6.2 EXPERIMENTAL DESIGN AND METHODOLOGY

A simple peg-insertion paradigm was used to quantify the effectiveness of overlaid audio localization percepts to enhance operator performance in a telemanipulation task. Seven human subjects were put through a battery of tests which included a number of training sessions followed by a data collection period. The overall testing procedure was broken up into three parts to be referred to as Test I, Test II, and Test III. Test I was a baseline test of basic human performance in which subjects performed the peg insertion task *in-person* while standing at the task-board. This was performed without the aid of any auditory perceptual overlays. Test II was a baseline test of telepresence performance in which subjects performed the peg insertion task through the exoskeleton system but again had no access to auditory perceptual overlays. Test III was the comparison test in which each of three audio fixture designs were implemented.

6.2.1 The Peg Insertion Task

The basic task required of subjects was to remove a peg from a designated start-hole and insert the peg in a designated target-hole as fast as possible. Data collected from such trials consisted of the time elapsed between the instant the peg left the start-hole until fully inserted in the target-hole. Although previous testing of two-dimensional audio fixtures used a formal Fitts' Law paradigm (as described in Chapter 2), the new round of testing implemented a modified peg-insertion paradigm to allow for a more efficient testing procedure. This modified Fitts' Law testing procedure reduced the trial count by using only one peg size (1.50 cm) rather than a traditional set of three peg sizes. Although a range of task difficulties

was not tested, this peg-insertion paradigm proved effective in previous three-dimensional fixture tests as described in Chapter 5.

Two peg holes of 2.0 cm diameter were used in this study which were equally spaced 4.0 cm from the center of the task board as shown in Figure 6.1. These holes were referred to as *Left* and *Right* as indicated on the figure by L and R respectively. The holes were closely spaced so as to only require peg motions of small amplitude. Small amplitude motions were chosen for this study to reduce the encumbrance effect imposed by the exoskeleton (resulting from the added inertia to the subjects' arm). It was found that these small amplitude motions could be performed for many trials without significant fatigue. Since a small amplitude motion was used, a tight tolerance between peg and hole was used to produce a high level of task difficulty. The 1.5 cm peg used in this experiment made for a very challenging insertion task when performing remotely though the telepresence system. In fact, three 30-minute practice sessions were required before all subjects became fully proficient.

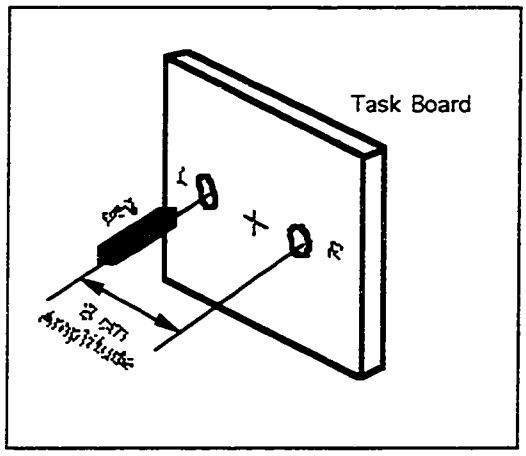

Figure 6.1: Peg Insertion Task

6.3.2 Testing Procedure:

Each of Test I, Test II, and Test III were divided into to a number of daily sessions which included a short warm-up followed by 30 recorded trials. Testing sessions only lasted between 10 to 20 minutes to minimize the fatigue effects.

Table 6.1: Overall Testing Schedule

TEST	SESSION	DESCRIPTION OF TRIALS
Test I		In Person Peg Insertion Trials: No Telepresence Baseline
	Session 1	practice trials
	Session 2	practice trials
	Session 3	practice trials
	Session 4	data trials
Test II		Telepresence Peg Insertion Trials: No Fixture Baseline
	Session 1	hardware familiarization
	Session 2	practice trials
	Session 3	practice trials
	Session 4	practice trials
	Session 5	data trials
Test III		Telepresence Peg Insertion Trials with Virtual Fixtures
*	Session 1	Hand Localizing Beacon Fixture: Linear
*	Session 2	Hole Localizing Beacon Fixture: Linear
*	Session 3	Hole Localizing Beacon Fixture: Non-Linear
	Session 4	no fixture baseline

* denotes that order of these trials were randomized

6.2.2.1 TEST I: In-Person Peg Insertion Trials

The experimental goal of Test I was to record a base-line of human performance for the peg-insertion task, free from the performance degradation resulting from remote presence. Ultimately, this data was recorded to compare the performance achieved in a telepresence task executed with virtual fixtures to the maximum performance achieved when performing the task in person. Thus it was desired to create an in-person task which was identical to the remote presence in every way possible except for the single distinction that the task was not performed through a telepresence interface. To account for the emcumbrance imposed by the exoskeleton during telepresence trials, the *Exoskeleton Encumbrance Simulator* (see Chapter 5 for details) was used during the in person trials.

The experimental procedure for Test I was as follows; Each subject was introduced to Task Board, Exoskeleton Encumbrance Simulator, and the peg-insertion task. They were instructed to perform the peg-

insertion tasks as quickly as possible (as long as they could consistently and confidently perform the task at that speed). Next they were asked to perform three practice sessions (one each day for three consecutive days). Each practice session began with 12 warm-up trials to allow subjects to get used to the feel of the hand grip and the insertion task. The warm-up was followed by 30 timed trials. At the end of the third practice session, results confirmed that learning had stabilized (i.e., variation among similar trials fell below 20% for all subjects tested). On the forth day, subjects were asked to perform a final test session which was used to collect the actual data for this experiment.

6.2.2.2 TEST II: Telepresence Peg Insertion Trials without Virtual Fixtures

After all of the subjects had performed Test I, they were introduced to the experimental hardware to be used in the telepresence peg-insertion trials. Subjects were placed in the operator space which included the FREFLEX exoskeleton and the stereo vision system. Subjects were taught how to operate the safety features for the exoskeleton system which included a hand switch and foot pedal which engaged the force reflection motors. Subjects were then introduced to remote workspace which included the MERLIN robot arm and the task-board.

Subjects were provided with an initial 30-minute session during which they were asked to familiarize themselves with the equipment. Many subjects needed this exposure period to simply get aquainted with the illusion of telepresence, to get used to the feel of the exoskeleton, and to get comfortable with controlling the remote manipulator. Once comfortable with the system, all subjects were asked to perform 3 practice sessions (one each day for three consecutive days) during which time they would become fully trained on the telemanipulation task. Once training was complete, subjects were asked to perform a final testing session which would serve as the actual data collection period.

6.2.2.3 TEST III: Telepresence Peg Insertion Trials with Virtual Fixtures

Once the subjects had gone through Test I and Test II, they were then exposed to the notion of *3D audio localization* virtual fixturing and the ALCS was introduced to them. Subjects were told that a set of three different audio virtual fixture designs would be introduced and that these fixture were intended to enhance their performance in the peg-insertion task. Although all subjects were asked to make an attempt at using these perceptual overlays to enhance their performance, no guidance was given as to how the fixtures should be used. During the three day testing period, subjects would perform a short warm up followed by 30 data collection trials on each testing day. A new virtual fixture design was introduced during

each of the three days. The order in which subjects were tested on the three different fixture designs was randomize to avoid ordering effects.

During Test III, subjects were only allowed a short 15 trial warm-up session with each virtual fixture design before performing the data-collection trials. There were a number of reasons for not extensively training subjects on the various fixture designs during Test III. First, all subjects had already been fully trained on the peg-insertion task (without fixtures) and were comfortable with using the hardware to perform peg-insertions. Thus there was less of a need for an extensive four day training procedure as was used in Test II. Second, the basic intent of virtual fixturing is to provide operators with generic overlays which can be implemented when confronted with a challenging task upon first encounter. Thus, taking performance measurements after extensive training in using a given fixture to perform a particular task would not reflect the intended application.

The final session of Test III was a baseline test in which subjects performed the task once again without the aid of auditory virtual fixtures. These baseline trials were recorded and compared with the data collected in the final session of Test II. Such a comparison was performed in order to quantify how much of the performance improvement occurring during Test III could be attributed to learning the basic task. As it turned out, subjects showed no significant performance improvement in the final baseline session when compared with performance in the final session of Test II. This suggests that subjects were fully trained by the end of Test II, and that the performance improvements observed during the fixture trials of Test III was not the result of a simple learning effect.

6.2.3 Post Testing Interviews

After the completion of each of the three audio virtual fixture sessions of Test III, subjects were given post testing interviews during which they were asked to described their impression of the various fixture designs. Subjects were not allowed to see the quantitative results until after interview sessions to avoid influencing their feedback. The intent of these interviews was to gain insight into why certain fixture designs may have been successful in enhancing operator performance.

6.2.4 The Virtual Fixture Models

The Audio Localization Cue Synthesizer (ALCS) was used to generate binaural perceptual representations of audio cues positioned spatially in a three-dimensional environment. The output of the ALCS was controlled by software running on a 486 microprocessor. These software routines were responsible for modeling the three virtual fixture designs displayed to the subjects through stereo

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

headphones. These auditory fixtures were perceived by subjects as being overlaid on top their perceptual representation of the remote telepresence environment. Although these audio fixtures were virtual constructs artificially introduced into the user's perceptual reality, subjects attributed these overlaid percepts to the physical reality of the remote environment. In other words: Although the virtual stimuli were displayed through headphones proximally, they were perceived as audio signals emanating from distant sources within the remote environment.

Each virtual fixture was modeled by a simple mathematical function which took as input the Cartesian position of the FREFLEX hand grip in three-dimensional space and returned a localization command as output to the ALCS. The three virtual fixture models implemented in this experiment were: Linear-Peg-Beacon, Linear-Hole-Beacon, and Nonlinear-Hole-Beacon. These particular fixture elements were *not* necessarily the most effective audio fixture designs for the given peg-insertion task. Rather, this set of virtual fixtures was chosen to simply reflect a variety of distinct perceptual overlays that make use of binaural audio cues. The following sections describe each of these virtual fixtures and their underlying mathematical models.

6.2.4.1 Linear-Peg-Beacon Fixture

The Linear-Peg-Beacon virtual fixture was composed of an audio signal whose perceived origin could be mapped through the ALCS to arbitrary 3D locations within the operator's perceptual reality. The particular mapping employed by this fixture design was simply a 20:1 linear scaling between the operator's hand position and the 3D origin of the auditory signal. In other words, the perceived localization of the auditory signal was mapped to follow the operator's hand position in 3D space, but was exaggerated such that when the operator moved his or her hand 5 cm to the right, the sound appeared to move 100 cm to the right. This perceptual overlay provided an audio localization cue, or *beacon*, indicating relative (not absolute) position of the operator's right hand (the hand which performed the peg insertion task) in 3D space.

Because subjects were immersed within the remote telepresence workspace when using this fixture, they did not attribute the audio localization cue to changes in *hand position*, but rather attributed it to changes in robot *end-effector* position. Thus they did not conceptualize this fixture as a kinesthetic beacon indicating hand position, but rather as an end-effector beacon indicating peg position within the remote site. This effect, unanimously confirmed by all subjects during post testing interviews, is a particularly insightful example of the *distal attribution* phenomenon discussed in Chapter 1 and Chapter 4. When a user is confronted with a proximal stimulus such that a lawful relation exists between afference and efference, the user accepts this stimulus as a genuine part of the ambient reality. Since the ambient reality

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

of this experiment was a remote telepresence worksite, users attributed this signal to that environment rather than the physical surroundings. Because the afferent feedback from this ambient reality contained a visual representation of the remote manipulator but did not contain a visual representation of hand position, the beacon was attributed to the end-effector rather than the hand. This observed phenomenon supports the basic premise on which this thesis is based, namely that *reality* is a *perceptual* rather than *physical* construct.

6.2.4.2 Linear-Hole-Beacon Fixture

The Linear-Hole-Beacon virtual fixture was composed of a the same audio signal used in the previous example, but implemented a different mapping though the ALCS. While the previous fixture mapped the beacon signal to the operator's hand position, this one mapped the perceived *observation point* to operator hand position while keeping the beacon signal at a fixed location in 3D space. The fixed location of the beacon signal corresponded with the 3D position of the target-hole in the remote environment. The net result was an abstract perceptual environment in which the operator always perceived the beacon signal as emanating from the target-hole, but was given the illusion that his or her observation point changing with hand position. Again this mapping was not attributed to hand position, but rather was attributed to the ambient reality (i.e., robot motion in the remote environment). Thus the user felt as though he or she was hearing a beacon emanating from the target-hole from the perspective of *riding on* the end-effectorof the remote manipulator arm. Thus, when the peg was aligned with the target, the beacon sounded as though it was directly ahead; when the peg was off to the right of the target hole, the beacon sounded as if it was off to the user's left. The mapping employed for this fixture design was essentially the *inverse* of the 20:1 linear mapping used in the previous Linear-Peg-Beacon fixture.

6.3.4.3 Nonlinear-Hole-Beacon Fixture

This virtual fixture design was composed of the same audio signal used in the previous two examples, but used a different ALCS mapping. The mapping was similar in concept to the Linear-Hole-Beacon in that the beacon sound was perceived as emanating from a fixed location at the target-hole. The difference was that while the previous fixture used a linear mapping between hand position and observation location, the current fixture design used a nonlinear mapping. The intent of this nonlinear function (shown in Figure 6.2) was to amplify the perceptual resolution of the localization cue as the user approached the target-hole. This mapping provided the greatest magnitude spatial change in the localizing cue during the peg insertion portion of the task when such cues were most beneficial. In other words, when the user's hand was far from the target, there was very little change in the audio signal; When the user's hand came near the target-hole, the change in audio localization spanned a large distance.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

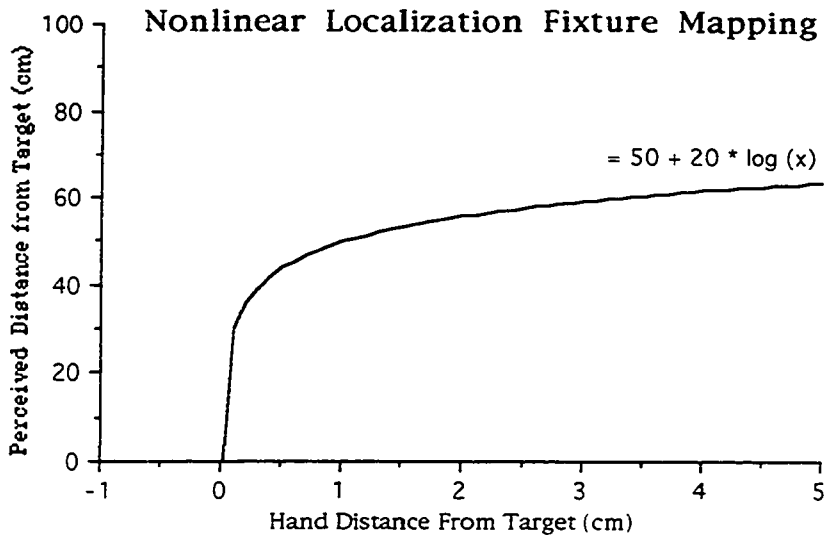

Figure 6.2: Nonlinear Auditory Localization Fixture

6.3 RESULTS

Movement time (mt), computed as the time elapsed during peg-insertion task execution, was recorded for each data-collection trial. Mean movement time and coefficient of variation (COV) values were computed for each subject performing the task using each of the fixture configurations (Linear-Hand-Beacon, the Linear-Hole-Beacon, and Nonlinear-Hole-Beacon). These values were also computed for the In-Person Baseline and the No-Fixture Baseline sessions. This data was further compiled by computing mean movement times and coefficient of variations across all 6 subjects as shown in Table 6.2.

Table 6.2: Movement Time Results, averaged across all six subjects

SESSION	MEAN mt	COV
In Person Baseline	586	.22
No Fixture Baseline	1785	.22
Linear Hand Beacon Fixture	1503	.24
Linear Hole Beacon Fixture	1547	.24
Nonlinear Hole Beacon Fixture	1419	.22

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

6.4 ANALYSIS

To best compare operator performance between trials of each virtual fixture configuration, the Index of Telepresence Performance (I.T.P.) was computed as follows:

I.T.P. = Movement Time in-person / Movement Time in Telepresence Task

As defined in Section 5.6, the ITP reflects telepresence performance as a percentage of the in-person baseline. An I.T.P. of 1.0 means that performance was equally good when using the telepresence system as in person. An I.T.P of less than 1.0 reflects a performance degradation while an I.T.P of more than 1.0 reflects a performance enhancement above natural ability. Table 6.3, below, shows the ITP for this set of tests, computed from the average times in Table 6.2.

Table 6.3: Index of Telepresence Performance (I.T.P.) for all Sessions:

	In Person Baseline	Telepresence No fixture	Linear Hand Beacon	Linear Hole Beacon	Nonlinear Hole Beacon
MEAN I.T.P.	1.00	0.34	0.40	0.39	0.42

A final step of data processing provided a direct comparison between the effects each virtual fixture has on operator performance. The *percentage performance increase* of I.T.P was computed for each virtual fixture design as compared to the no-fixture baseline case. These values are shown below in Table 6.4

Table 6.4: Percent Increase In I.T.P. For Each Virtual Fixture Session

	Linear Hand Beacon	Linear Hole Beacon	Nonlinear Hole Beacon
% Increase	**18.6** % ± 7.8	**15.2** % ± 6.3	**25.6** % ± 12.2

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

6.5 DISCUSSION & CONCLUSIONS

6.5.1 In-Person Baseline Trials Compared to No-Fixture Baseline Trials

What may be the most striking observation implied by the above results is the dramatic difference in operator performance between the in-person baseline trials and the no-fixture telepresence baseline trials. When performing the task in person, peg-insertion times were on average 586 ms but when using the telepresence link to the remote site, mean movement times grew to 1785 ms. This huge drop in performance is similar to what was observed in the previous set of experiments described in Chapter 5 in which mean movement time jumped from 625 ms to 2405 ms between baseline sessions. It should be reminded that this performance drop can *not* be easily attributed to the added physical encumbrance imposed by the telepresence hardware because all in-person trials were performed using the *Exoskeleton Encumbrance Simulator* which matched the physical constraints imposed by the exoskeleton. As described in detail in Chapter 5, a likely explanation for the dramatic drop in operator performance in the telepresence baseline trials as compared to the in-person baseline trials is a degradation in the illusion of operator *presence* at the worksite. Such a drop in operator presence was likely the result of reduced modalities through which afference and efference could be lawfully related.

6.5.2 Virtual Fixture Trials Compared to No-Fixture Baseline Trials

Turning next to the peg-insertion trials performed with the aid of the various auditory virtual fixture designs, we find that all three binaural perceptual overlays resulted in measurable performance improvements as compared to the no-fixture baseline. Performance improvements of 19%, 15% and 26% were recorded for the Linear-Hand-Beacon, Linear-Hole-Beacon, and Nonlinear-Hole-Beacon virtual fixture designs respectively. One possible explanation for the increased performance is that an additional modality was provided through which afference could be related to efference (in this case kinesthetic commands could be related to binaural auditory localization cues). Such an additional relation between afference and efference may enhance operator presence in the remote site and ultimately enhance task performance. The design and application of binaural audio percepts can be seen as a controlled enhancement of the operator's perceptual reality with the intent of improving operator performance by increasing the illusion of presence in the remote site. Although such conclusions are consistent with current theories about the perceptual requirements for operator presence, post testing interviews did not support these ideas. Subjects did not report having an improved sensation of *presence* in the remote site as a result of the overlaid binaural auditory percepts. Of course, the fact that subjects did not verbalize enhanced presence does not necessarily imply that such a phenomenon did not occur. Nevertheless, explanations other than enhanced presence should be considered in light of subject feedback.

Rather than using the perceptual overlays to feel more *present* in the remote site, subjects reported using the auditory perceptual overlays to acquire a better sense of peg *position* and peg *velocity* with respect to the remote environment. Although enhanced spatial localization in the remote worksite can be seen as a ramification of increased presence, for the sake of simplicity *enhanced spatial localization* is discussed as a unique phenomenon. All subjects reported using the binaural auditory cues as a direct indication of the 3D location of the peg within the remote environment. Some subjects also reported using the rate of change in binaural auditory cues as a direct indication of the 3D velocity of the peg with respect to the remote site. This usage implies that the overlaid audio information was *attributed* to the remote site and perceived as a genuine property of the remote task. Psychoacoustic researchers generally refer to this remote attribution of a binaural cue as the *externalization of an auditory image* [Durlach 1992].

The fact that subjects attributed the binaural audio percepts to physical properties of the remote peg (such as its position and velocity) is a particularly interesting example of the *distal attribution* when we consider that the mappings between proximal and distal stimuli were very different from what physical *realism* would predict. A physically *realistic* sound emanating from a distant source has a particular mapping between spatial position and binaural audio representation. Recall that the two linear virtual fixtures employed a 20:1 exaggeration of this mapping, while the third, nonlinear virtual fixture used a logarithmic distortion of the mapping. Despite the fact that the given perceptual overlays were not consistent with what physical accuracy would predict, these virtual embellishments of the operator's *perceptual reality* provided useful and believable localizing references to the remote site. So much so that all subjects were able to use these abstract perceptual cues to enhance performance in a rigidly structured peg-insertion task. In fact, the virtual fixture design which deviated most from physical realism, namely the logarithmic binaural mapping, happened to be the fixture which most enhanced operator performance. This is likely because the logarithmic mapping provided exaggerated localization cues during theactual insertion portion of the task, when localization information was most critical.

These results of this study provide a strong confirmation of the basic premise on which this thesis is founded, namely that generating believable virtual percepts (i.e. percepts which users will accept as part of their ambient reality) is not dependent upon accurate physical modeling. Rather, the generation of realistic and useful virtual percepts depends upon the designer's ability to provide a perceptual reality in which lawful relations exist between afference and efference. Such lawful relations are not restricted to physical realism. Abstract relations such as exaggerated linear mappings and logarithmic mappings were found adequate to provide believable and effective virtual percepts.

In addition to providing enhanced localization to the remote environment, the binaural virtual fixture designs could enhance performance in other ways. Many subjects reported using the overlaid perceptual cues as an additional means by which to gauge task performance. For example, subjects reported being able to assess an efficient peg trajectory as not just looking and feeling a certain way, but as also spatially *sounding* a particular way. Thus subjects could possibly use audio virtual fixtures to develop more complete or efficient *mental models* of task execution.

Chapter Seven

"The designer must develop a conceptual model that is appropriate for the user, that captures the important part of the operation... "

-Donald Norman
The Design of Everyday Things

---

7. The Perceptual Workstation

Virtual Fixtures are computer generated percepts designed to enhance operator performance in manual tasks. The use of well designed virtual fixtures have been shown in earlier chapters of this document to be very effective in both direct and time delayed telemanipulation activities. Simple fixtures composed of rigid haptic surfaces and auditory fields have been shown to increase operator performance up to 70% in a standardized peg-insertion task. Simple haptic fixtures composed of intersecting rigid planes have also been shown to reduce the performance degradation due to time delays by up to 36% and 44% for delays of 250ms and 450ms respectively. The use of more abstract haptic fixtures such as snap-to-planes and snap-to-lines have been shown to enhance operator performance as much as 102% in 3D telemanipulation tasks. The use of binaural audio localization fixtures were shown to enhance performance as much as 25%. In general, subject testing of the virtual fixturing concept has revealed that users are comfortable with the *virtual fixture metaphor* across a wide range of fixture designs and have a clear conceptual model of how these abstract perceptual overlays can be used effectively to assist performance in manual tasks.

Although previous chapters have clearly demonstrated that virtual fixturing is effective when subjects are provided with pre-defined fixtures, issues related to the *development* of virtual fixture designs have not yet been addressed. The development effort required to construct effective fixtures has great bearing upon the potential application and flexibility of virtual fixturing. If the generation of effective virtual fixtures requires a substantial design effort, the use of virtual fixturing may not be worth the development

time for certain applications. Thus far, empirical testing has shown that relatively simple fixtures designs can produce substantial performance increases. This is a strong indication that a significant design effort may *not be required* to construct effective perceptual overlays. Nonetheless, all fixtures tested thus far have been pre-defined by the experimenter who had previous knowledge regarding the configuration of the task and the layout of the remote workspace. If the generation of effective fixtures requires intimate prior knowledge of the remote workspace and remote task, the up-front effort that would be put into designing a fixture for a given task might be better spent simply programming the remote manipulator to perform the well-known task autonomously. If, on the other hand, effective virtual fixtures could be rapidly developed by the human operator upon first encountering a task, virtual fixtures could be used in unstructured or changing environments unsuitable for autonomous systems. If we consider that the basic goal of a telepresence system is to extend an operator's natural manual and cognitive abilities to a remote environment, it seems that providing operators with the ability to interactively design and modify assistive virtual fixtures may be essential to preserving the flexibility of the human in the loop.

Because the ability for interactive fixture design is important to the effectiveness and flexibility of the virtual fixture concept, the development of a *Perceptual Workstation* was proposed in Chapter 1. It has been envisioned as an interactive development environment intended to enable operators to design and implement their own virtual fixtures when confronted with a given task. The basic premise is as follows: If an operator is provided with a *toolbox* of simple perceptual building blocks, he or she can construct custom tailored virtual fixtures that suit the task at hand. With a few simple commands issued to the Perceptual Workstation, a user could instruct the system to overlay a particular arrangement of abstract perceptual information on top of the reflected sensory feedback from the remote worksite. Similar to the user who grabs a ruler and lays it down on a page when confronted with a straight line drawing task, an operator using the Perceptual Workstation, could reach into the given toolbox of abstract perceptual elements and call upon equally generic overlays to enhance performance in complex three-dimensional tasks.

To facilitate the design and construction of virtual fixtures, it is proposed that the given toolbox of perceptual elements be composed of easy-to-understand yet versatile-to-use perceptual information. These basic elements might consist of simple haptic percepts such as rigid walls and compliant surfaces. The toolbox might also include more abstract sensations such as viscous liquid regions or attractive and repulsive fields. The toolbox might include auditory percepts such as spatial fields or binaural localizing cues. The toolbox might even include tactile percepts such as vibrations and textures. In addition, all fixtures might be given a visual representation to add to perceptual channels listed above. The intent of such a toolbox would be to provide users with a rapid means of accessing perceptual elements which can be used alone or in simple combinations to construct effective virtual fixtures for a given task.

In preceding chapters, it has been demonstrated empirically that the virtual fixture metaphor provides users with a conceptual model which allows them to *make use of* abstract perceptual overlays to enhance their performance in given tasks. It has not yet been shown, however, that the virtual fixture metaphor is powerful enough to allow users to actually *design their own* perceptual overlays. The virtual fixture metaphor does make the notion of perceptual overlays more manageable, but the idea of designing perceptual overlays composed of abstract sensory information is a concept that typical users may not be even remotely familiar with. Thus, the fixture metaphor may not provide enough of a conceptual framework to facilitate the design process. To evaluate if naive users could use the proposed metaphor, design their own perceptual overlays, and enhance their performance in a given task, an empirical study was performed. This study, to be described throughout the rest of this chapter, provided naive test subjects with a prototype of the proposed Perceptual Workstation. Once introduced to the workstation, subjects were asked to design their own virtual fixtures for a given telemanipulation task. The goal of the study was to quantify whether naive subjects could use the Perceptual Workstation environment and design effective virtual fixtures from a toolbox of abstract perceptual elements.

7.1 EXPERIMENTAL HARDWARE & SOFTWARE

7.1.1 Experimental Hardware

A limited prototype of a Perceptual Workstation was developed using an Intel-486 computer interfaced to a two degree-of-freedom (dof) force reflecting joystick [Adelstein, 1989; Adelstein, 1991]. Each dof of the joystick was powered by disk armature permanent magnet motors and instrumented with optical encoders to sense position, tachometers for velocity, accelerometers, and a force transducer. The motors could produce continuous forces up to a maximum of 20N with zero cogging and negligible force ripple from DC up to 58 Hz at the joystick handle. The minimum friction force threshold of the uncompensated joystick was 1N. The joystick handle's passive inertia corresponded to a mass of 0.35 kg at the hand. The joystick was operated in these experiments under purely digital control through an A/D and D/A card with DMA control on an ISA bus Intel 486DX-50 based personal computer. The controller update rates for the haptic models simulated in these experiments exceeded 2 kHz.

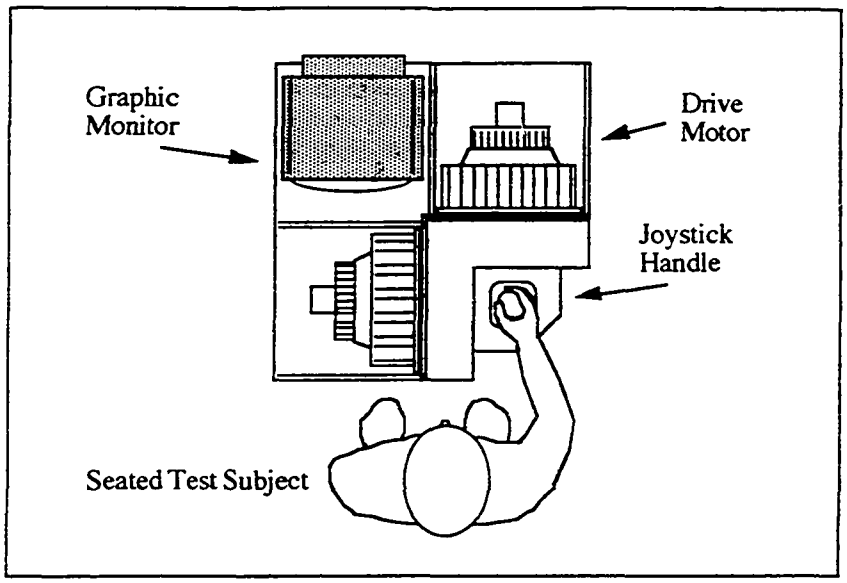

Figure 7.1: Experimental Test Setup

Subjects sat facing the joystick and the computer monitor as depicted in Figure 7.1. The handle, which was at a height of 1 m from the floor, was grasped in the right hand. Virtual fixtures were displayed visually on the monitor and haptically on the force reflecting joystick. The relative position of the joystick and the monitor allowed for a natural mapping between kinesthetic motions of the hand and graphic output on the screen.

7.1.2 Experimental Software System: *Percept-Station*

7.1.2.1 Background

An experimental software package was developed to demonstrate the feasibility of the Perceptual Workstation concept. This software package, referred to as *Percept-Station*, was designed to provide operators with a toolbox of perceptual elements which could be overlaid on top of a two-dimensional remote environment. Percept-Station allowed for the interactive design and implementation of virtual fixtures composed of the following haptic elements:

-Rigid Horizontal Walls
-Rigid Vertical Walls
-Viscous-liquid-regions
-Horizontal Grating Textures
-Vertical Grating Textures
-Repulsive Fields
-Bumps

This particular set of elements was chosen to provide a diverse array of haptic sensations for fixture design. The intent was *not* to test an optimal set of perceptual elements for an ideal perceptual workstation, but rather to simply test whether subjects could make use of a given toolbox of perceptual elements to design coherent and effective perceptual overlays.

7.1.2.2: The Remote Environment: A Virtual Model

In order to simplify the hardware requirements for this experiment, Percept- Station implemented a *virtual model* of the remote environment rather than using actual sensory feedback reflected from a real remote site. The use of a virtual environment for the telepresence task eliminated the need for a real time camera hookup to the remote worksite while also eliminating the need for an actual remote manipulator. A virtual remote environment was also used so that subjects could err while performing the task without severe consequences. It was felt that using a computer simulation of the remote environment, rather than an actual telerobotic link, provided a significant savings in equipment, complexity, and development time while still providing an adequate testbed for assessing the feasibility of the perceptual workstation concept. Although the remote environment and remote task were graphical simulations, all subjects were presented with a realistic telemanipulation scenario before performing the tests in order to motivate their activities in terms of a real world paradigm.

Since the Percept-Station software package was designed to support two-dimensional manipulation tasks, a planar remote environment was chosen for this study. To provide a realistic and motivational paradigm in two dimensions, the simulated remote environments were represented as top-view floor plans of a fictional remote location. A remote manipulator was represented in this remote environment as a hexagonal graphical icon whose position the user could control using the force reflecting joystick. Users were presented with a scenario that they were controlling a distant mobile robot (represented by the graphical icon) which could be moved around inside of a remote laboratory (represented by the graphical floor plan). Although an abstract representation, the scenario provided an adequate testbed for remote manipulation tasks requiring high dexterity. For these tests, the virtual remote manipulator responded ideally to the users commands with no dynamic lag or communication delays. Although those types of distortions would likely be present in a true telemanipulation task, previous testing of virtual fixtures with time delay has confirmed that virtual fixturing remains effective even when time delays are present.

7.1.2.3 Fixture Design

Subjects were provided with a toolbox of perceptual elements and a graphical user interface with which to construct virtual fixtures. The constructed fixtures were overlaid on top of the representation of

the remote site. This means that the fixtures were displayed graphically and haptically in the same workspace as the perceptual information reflected from the remote environment. Thus the virtual fixtures were overlaid directly on top of the graphical floor-plan which represented the remote environment. Any actions made by the user within the workspace would interact with both the remote environment and the overlaid virtual fixtures as if both occupied the same space.

7.1.2.4 Graphical User Interface for Fixture Design

Percept-Station allowed subjects to design and implement virtual fixtures through a familiar graphical user interface. Fixtures were constructed and positioned with respect to the remote environment by simply drawing graphical representations of the fixture elements in the locations where they were to be implemented. For example, a repulsive field could be positioned at a particular location in the remote environment by simply moving the cursor to the desired position and commanding a repulsive field with an appropriate key stroke. A graphical representation of the repulsive field would then appear overlaid on top of the remote environment in the location where it would be haptically perceived. Users designed all virtual fixtures by first choosing from among a menu of possible perceptual elements and then drawing a graphic representation of the percept at the desired location in the remote environment. The drawing of virtual fixtures was achieved by manipulating a cross-hair cursor using arrows on the keypad. The interface was similar to many commercial drawing applications in which a cursor is controlled to draw and erase lines and other graphical elements. All subjects reported finding the user interface familiar in structure and easy to learn.

Each haptic percept in the toolbox had a unique visual representation which allowed the fixture elements to be described graphically on the layout of the remote environment. Each haptic percept also had anassociated keystroke which allowed the user to select it from among the other elements in the toolbox. For example, a user could describe haptic percepts composed of horizontal and vertical rigid walls by simply drawing a closed polygon of horizontal and vertical lines using the cursor. When the polygon was complete, the user would position the cursor inside of the region and then press <W> for RIGID WALL. When <W> was pressed, the drawn figure was filled-in and appeared as a solid polygon colored cyan. The edges of the polygon represented rigid walls that could not be penetrated when encountered by the force reflecting joystick. Similarly a user could describe haptic percepts composed of viscous-liquid-regions by drawing a closed polygon and pressing <L> for LIQUID REGION when inside the closed figure. When <L> was pressed, the figure was filled in and appeared as a solid green polygon. When the force reflecting joystick entered the viscous-liquid-region, the user would feel velocity proportional resistance as if dragging through a thick viscous liquid.

To generate a REPULSIVE FIELD emanating from a point source, the user simply moved the cursor to the desired emanation point and pressed <R> or <r> to create a strong or weak repulsive field respectively. When <R> or <r> was pressed, a graphical representation of concentric yellow rings was drawn with a size proportional to the field strength. When the force reflecting joystick approached such field sources, the user's hand would be repelled with a force inversely proportional to the square of the distance from the source.

To generate a TEXTURE GRATING, the user simply outlined the region to be textured using the cross hair cursor. When a closed figure was drawn, the user simply entered the closed figure and pressed <T> for texture. The user could either press <H> or <h> to indicated a horizontal texture grating of wide or narrow spacing or could press <V> or <v> to indicate a vertical texture grating of wide or narrow spacing. Once the grating direction and density was specified, green grid lines were drawn on top of the remote workspace to represent the desired texture in the appropriate region. When the force reflecting joystick encountered a textured region, the user would feel a sensation similar to dragging a stick across a grated surface.

To generate a BUMP overlaid on top of the remote workspace, the user simply pressed <B> for bump and used the cross hair cursor to specify the outline of the bump. A blue line was then drawn along the trajectory specified by the cross hairs. When the force reflecting joystick encountered the blue line, a bump was felt. Many subjects have compared the sensation of moving the joystick through a bump fixture with driving a car over a speed bump. The bump is a noticeable sensation but still easy to pass over.

The following table lists all of the perceptual elements included in the Percept-Station toolbox as well as the keystrokes required to specify the percept.

Table 7.1: Perceptual Elements Supported by the Percept-Station Toolbox

Haptic Percept	Keystroke
Rigid Wall	W
Liquid Region	L
Repulsive Field	R , r
Texture Grating	T: (H , h , V , v)
Bump	B

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

7.1.3 Haptic Sensation Algorithms

7.1.3.1 General Strategy to Haptic Sensation Generation

A number of computational algorithms were used to implement the haptic sensations produced by the Percept-Station toolbox. In all cases, haptic sensations were presented by generating forces as a function of joystick handle position alone or joystick handle position and velocity according to particular dynamic equations. These equations switched *on* and *off* when the joystick encountered particular fixture elements. For example, when the joystick encountered a liquid-region fixture element, a force was generated at the handle which was proportional to joystick velocity in the direction opposite to joystick motion. When the joystick was no longer within the liquid region, the force was no longer applied. Recall that a graphic representation of the remote manipulator slaved to joystick motion was also displayed in the remote environment. Thus, when watching the remote manipulator on the graphic display of the remote environment, the user would feel the onset of the haptic sensation at the instant the visual display showed the remote manipulator encountering the fixture. This coupling between visual feedback, haptic feedback, and kinesthetic activity, made for convincing virtual sensations.

The following sections describe in detail each of the haptic elements supported by Percept-Station and discusses the mathematical model used to generated it.

7.1.3.2 Horizontal and Vertical Rigid Walls

A virtual rigid wall is a computer generated haptic sensation similar to what a user would feel when pushing the joystick handle up against a physical hard-stop. Percept-Station allows users to interactively design and then interact with complex two dimensional virtual fixtures composed of any number of virtual rigid walls positioned anywhere in the workspace. As a consequence of this flexibility, the resultant computational burden restricted the software from supporting anything other than horizontally and vertically oriented walls. Although the experimental software package only supported horizontal and vertical walls, this should not imply that diagonal or curved surfaces cannot be included in a fully functional perceptual workstation.

The mathematical model representing the haptic rigid wall was derived through rigorous subject testing and perceptual modeling techniques described in Chapter 4. The particular model used here was a parallel combination of a linear spring element k=4000 N/m and a directional damper element with

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

b=100 N/(m/s). The simple mathematical model which generated the haptic sensations for the horizontal and vertical walls can be described as follows:

$$F=kD+bV$$

Where D represented the distance that the joystick penetrated the virtual wall in a direction normal to the surface and V represented the directional velocity of the joystick handle normal to the surface. Because applied forces only opposed motion normal to the virtual surface, the haptic wall felt perfectly smooth and frictionless when sliding tangent to the surface. The wall model is shown schematically in Figure 7.2 where the shaded circle represents the joystick's handle position and the dark vertical line represents the location of the virtual wall model.

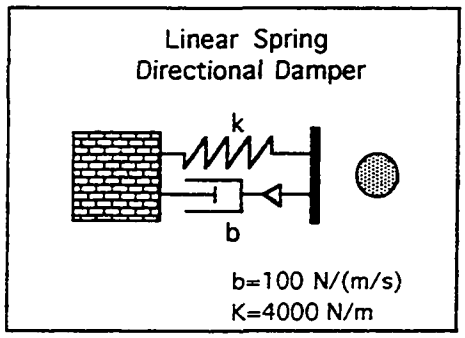

Figure 7.2: Virtual Wall Model: Linear Spring and Directional Damper in Parallel

Although the haptic virtual wall model was a simple parallel combination of a linear spring and directional damper as described above, the dynamic act of contacting the virtual wall required additional processing. An algorithm was needed to turn-on the haptic model when the joystick encountered the virtual wall and turn-off the haptic model when the joystick handle pulled away from the wall. To maintain a stable transition between turning walls on and off, a thin barrier zone of *artificial hysterisis* was included. The hysterisis was such that the wall model turned on and off at slightly different locations. This eliminated vibrations resulting from the joystick sitting on the edge between contacting and not contacting the wall model. Although the barrier zone was only two encoder counts wide, all such vibrations were completely eliminated.

7.1.3.3 Viscous-liquid-regions

A virtual viscous-liquid-region is a computer generated haptic sensation similar to what would be felt when dragging a stick through a thick liquid. The mathematical model which described this haptic percept was a strong viscous damper in which the resultant force always acted to oppose the direction of motion of the joystick handle. Such a mathematical model can be represented graphically as shown in Figure 7.3. The magnitude of the resultant force was proportional to the velocity of the joystick handle through a damping constant b=100 N/(m/s). The haptic display model was switched on when the joystick handle entered a viscous-liquid-region and then switched off when the handle exited the region.

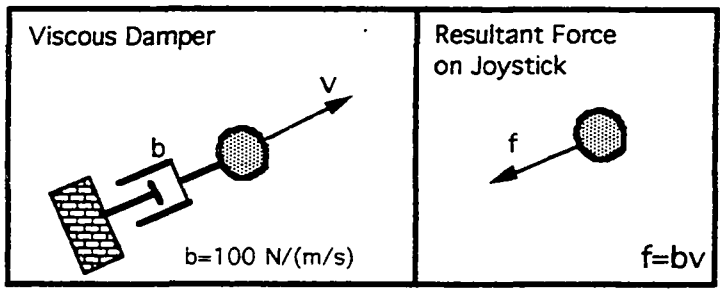

Figure 7.3: Haptic Model of Viscous-liquid-region: Viscous Damper

7.1.3.4 Repulsive Fields

To produce a haptic representation of a repulsive field emanating from a point source, a force was generated at the joystick handle proportional to the inverse square of the distance between the joystick and the designated point source. The direction of the repulsive force always coincided with a line drawn between the joystick handle and the point source as shown in Figure 7.4. The resulting sensation was a repellent force which was not perceivable throughout most of the workspace, but ramped up quickly as the joystick approached the source. Since the direction of the repulsive force was always opposite to the direction of the point source, when the joystick moved towards the repulsive field it was smoothly diverted around the source. Since Percept-Station only included the ability to generate horizontal and vertical walls but not surfaces of arbitrary topology, the addition of repulsive fields to the toolbox of perceptual elements allowed for interaction with smooth, rounded haptic sensations.

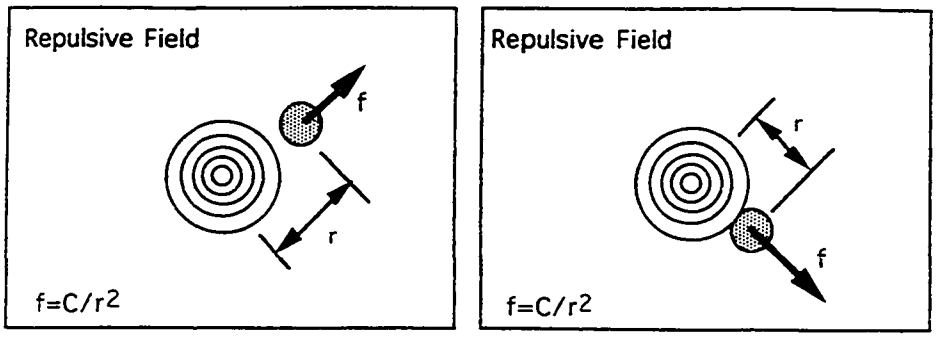

Figure 7.4: Point Source Repulsive Field Model

The construction of virtual fixtures from combinations of repulsive fields was particularly interesting because multiple repulsive fields applied in close proximity could influence each other. Neighboring fields could have summing or negating effects depending upon whether the joystick passed on the same side of multiple fields or in-between multiple fields. This allowed for the generation of very interesting virtual fixtures. For example, when two repulsive fields were placed side by side with a small gap between them as shown in Figure 7.5, the joystick handle would be guided directly down the middle between the two fields because lateral forces would balance but the forward backward forces would not. Users reported feeling like they snapped through a guided opening and shot out the other side.

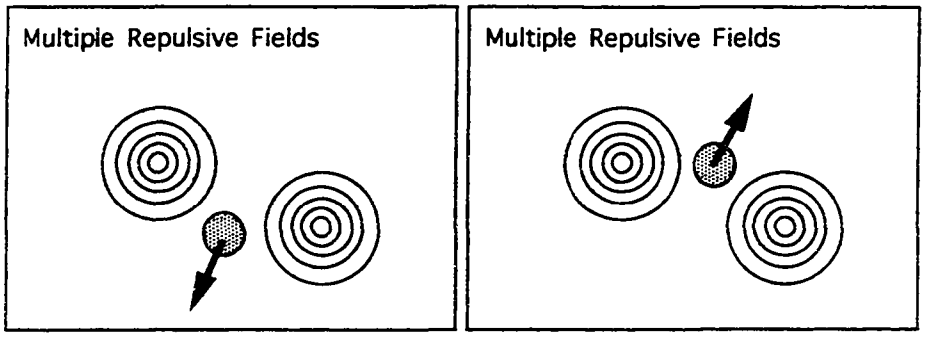

Figure 7.5: Implementation of Multiple Repulsive Fields to Compose Complex Sensations

7.1.3.5 Bumps

To represent the haptic sensation of encountering a small bump on a smooth surface, a perceptual model was required that would produce a noticeable sensation of contact but not disrupt motion. To create such a sensation, a simple model was generated which implemented similar features of

the virtual surface but which allowed the user to "break through" the surface when a minimal threshold was exceeded. Such a haptic model can be represented schematically as a spring damper system as shown in Figure 7.6 such that when a threshold displacement was exceeded, the opposing force dropped to zero and the joystick broke free from the physical resistance. If the displacement threshold was very small, the sensation was described by all subjects as a subtle bump or valley.

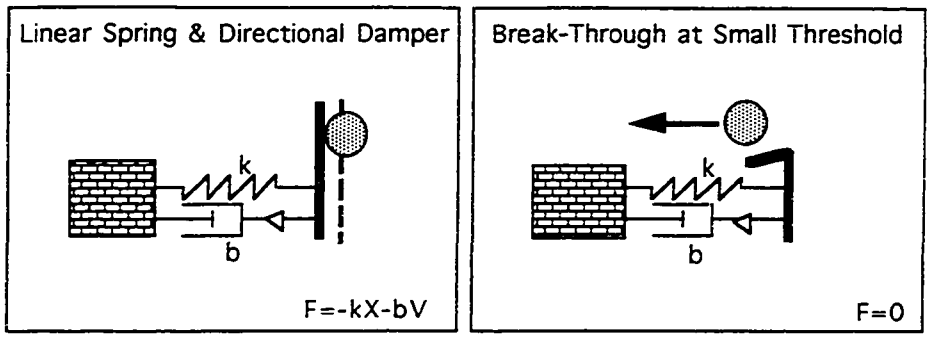

Figure 7.6: Two Phases of a Virtual Bump Model

7.1.3.6 Grating Textures

To produce a haptic sensation which represented sliding over a textured surface, a perceptual model was required to impart distinct textural vibrations in the joystick handle. This was generated with a virtual model consisting of a series of closely spaced virtual bumps, each of which was modeled as described previously in section 7.2.3.5. For these experiments, all textures were linearly aligned along the horizontal and vertical axis of the workspace. The user could specify any shaped region in the workspace and 'fill' the region with one of four possible haptic textures. When asked to describe the haptic sensations imposed by these virtual models, all subjects reported feeling like the joystick was being dragged over a textured region similar to sliding a stick across a metal grating. The four possible grating textures were composed of linear bumps arranged and oriented as shown below in Figure 7.7.

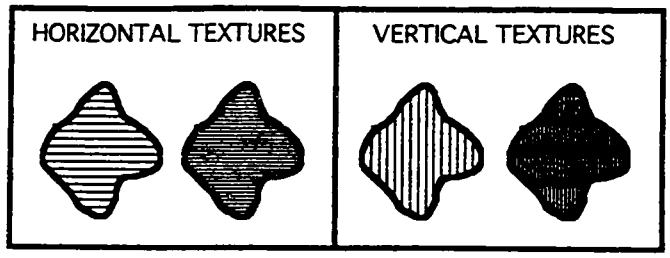

Figure 7.7: Four Configurations of Virtual Textures Supported by Percept-Station

7.2. EXPERIMENTAL DESIGN AND METHODOLOGY

7.2.1 Experiment Overview and Testing Goals

A simulated telemanipulation task was developed to quantify the effect of virtual fixtures upon human performance. While previous studies have investigated the effect of virtual fixtures in simple peg-insertion trials, the present study was intended to explore of the effectiveness of virtual fixturing in a more complex and appropriate telemanipulation scenario. It was desired to test human performance in a task which required a high degree of manual dexterity, continued concentration and mental planning during execution, and constant visual scanning of the workspace for efficient task completion. Although such a task would not provide as clean and controlled a test-bed for the evaluation of human performance as did the standardized peg-insertion trials, it was believed that a more realistic and challenging paradigm was required to demonstrate the effectiveness of the virtual fixturing.

A remote hazardous clean-up task was chosen as an idealized paradigm for the present set of tests because of the richness of task complexity as well as the appropriateness of such a task to real-world telemanipulation applications. The task required subjects maneuver the virtual telemanipulator through the remote workspace, pick up targets that were specified as *hazardous material*, and return such targets to the designated clean-up receptacle located in the corner of the workspace. Because there were numerous obstacles scattered throughout the workspace which had to be avoided, the task required a high level of manual dexterity, hand-eye coordination, and constant concentration. Because there were multiple targets scattered throughout the workspace, efficient execution of the task required subjects to visually scan the remote environment to locate future targets and plan an efficient path. This resulted in a task which was highly taxing upon the visual modality because subjects had to devote intense visual attention to guiding the remote manipulator past obstacles while at the same time scanning for new targets and planning effective paths.

In addition to simply demonstrating the effectiveness of virtual fixturing in a realistic telemanipulation task, an important goal of these experiments was to demonstrate that naive operators, with minimal training, could design and use their own virtual fixtures and enhance their own performance in a given task. While previous studies have demonstrated that virtual fixtures can greatly enhance performance, all previous tests provided the subjects with a specific set of virtual fixtures. Using the Percept-Station software package, this study was intended to be the first exploration of virtual fixturing in which operators would be provided with only a specific task and be required to design and implement their own virtual fixtures.

Subjects were given instruction for proper execution of the remote clean-up task and on how to use the Percept-Station software interface to design virtual fixtures. Subjects were then asked to design, construct, and implement their own virtual fixtures to assist their performance in the remote hazardous cleanup task. Operator performance was quantified through a simple scoring algorithm in which one point was awarded for each target collected in the given time period. One point was subtracted for each collision with obstacles in the workspace. Subjects were instructed to maximize their score during task execution. Post testing interviews were conducted to gain further insights into the effectiveness of virtual fixturing.

7.2.2 Hazardous Cleanup Task

As noted above, subjects were presented with a hazardous cleanup telemanipulation paradigm. The remote environment was a top view floor plan of a distant laboratory and the telemanipulator was represented as small graphical icon which could be maneuvered through the laboratory by using the force reflecting joystick interface. Scattered throughout the workspace were 6 targets to be collected and returned to a specified receptacle in the corner of the workspace. The floor plan itself was composed of a patchwork of shaded regions which specified obstacles to be avoided by the telemanipulator. The scenario was that the shaded regions represented dangerous, volatile, or fragile objects in the workspace that should not be bumped by the remote manipulator. Thus great care was required to maneuver the remote manipulator through the workspace without contacting the obstacle regions.

The targets were represented as small red circles at various locations in the workspace. Subjects were instructed to maneuver the telemanipulator throughout the floor plan and collect the designated targets as fast as they could. The challenge was to collect targets quickly while carefully avoiding contact with the specified obstacle regions. To require maximum travel throughout the workspace, and thus create a highly taxing testing scenario, subjects were instructed that the remote manipulator could only carry a single target at a time. Thus the remote manipulator had to be maneuvered to a specific target, pick up the target, and return that target to the clean-up receptacle before the next target could be retrieved. To pick up a target,

subjects simply needed to touch the target with the remote manipulator. When the graphical icon for the telerobot encountered a target, the target would disappear and the robot icon would change color to specify that the target had been retrieved and could now be transported. To drop off the target in the receptacle, the subject simply needed to maneuver the robot into the receptacle region. When the robot contacted the receptacle, the target would automatically be dropped and the robot icon would return to its original color to indicate that it was free to collect another target.

Subjects were instructed to collect as many targets as they could in a given time period of 90 seconds. The initial display of the remote workspace contained 6 targets. Once all six targets on the initial screen had been collected, a new set of six targets appeared in the same locations as the original targets. Thus subjects could continually collect as many targets as possible within the 90 second time period, the set of targets being refreshed after each screen was cleared. Because each new screen of targets was identical to the previous, a single virtual fixture could be designed for the initial layout of the workspace and be used for the entire task. This was done to allow a single virtual fixture to be designed for a task requiring the clean-up of many targets without having all the targets clutter the screen at once.

To quantify performance in the remote cleanup task, a simple algorithm was used to compute a score for each test session. Every time a target was picked up and returned to the receptacle, a point was awarded to the subject's score. Every time the remote manipulator contacted an obstacle, a point was subtracted from the subject's score. If the subject was carrying a target back to the clean-up receptacle at the time that contact with the obstacle occurred, a point was lost and the target being carried was lost so a new one had to be retrieved. Thus to achieve a high score in the testing paradigm, the subjects not only had to perform the task as fast as they could in the 90 seconds given, but they had to be highly accurate in avoiding obstacles because the penalty for collisions was severe.

A sample workspace is shown below in Figure 7.8. The shaded regions represent the obstacles to be avoided, the small circles represent the targets to be collected, and the cross-hatched region represents the receptacle to which each target is brought. The hexagonal icon represents the location of the remote manipulator in the workspace.

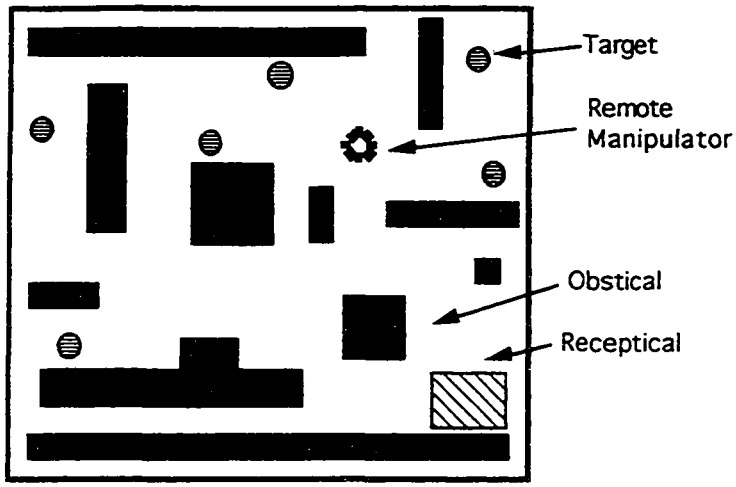

Figure 7.8. Remote Workspace for Cleanup Task I.

7.2.3 Instructions Given To Test Subjects

Subjects were all given an identical sheet to read which presented the testing scenario for the remote cleanup task. The goal was to motivate the abstract task with a realistic paradigm for telemanipulation. The text of these instructions are listed below as presented to the subjects:

THE SCENARIO: Imagine that you are looking down on the floor plan of a distant laboratory and that the joystick in front of you controls the motion of a mobile robot in that remote lab. Notice how the joystick controls the position of the small green icon on the screen. This icon represents the position of this remote robot. Now imagine that the lab has had an accident that resulted in the release of hazardous material. Your job is to control this remote robot from the safety of your distant control center to clean up the hazardous waste in the Lab. The small red dots scattered around the lab represents the hazardous waste that you must clean up.

THE TASK: You are to collect the red objects by touching them with the remote robot and then carrying them to the green cleanup receptacle in the corner of the workspace. The robot has been designed so that it can only carry a single object at a time so you must make a separate trip for each. After all six objects on the screen have been collected, a new screen will appear with six more objects to collect. The goal is to collect as many objects as you can in a 90 second time period.

THE CHALLENGE: Although the goal is to collect as many objects as you can in the 90 second period provided, you must be careful not to drive the remote robot into the gray shaded regions on the screen. These regions represent dangerous or fragile obstacles in the remote environment that must be avoided at all costs. Driving the robot into the gray region represents a destructive failure which would further contaminate the laboratory.

SCORING: You will get 1 point for every object that is collected and returned to the green cleanup receptacle You will loose one point for every time you enter a gray hazardous region. If you enter a gray region, you will also drop the object you may have been carrying at the time in addition to having your score reduced by one. The goal is to get as high a score as you can in the time given.

7.2.4 Testing Protocol

7.2.4.1 Testing Sessions

The overall testing procedure was broken up into two sessions: Session I and Session II. Session I lasted from 60 to 90 minutes in which subjects were familiarized with the experimental hardware, trained in the remote cleanup task, trained in how to use the Percept-Station interface to design virtual fixtures, and tested in a number of trials using a variety virtual fixture elements. During this initial testing session, all trials were performed using the same remote workspace (i.e., all workspaces had the same layout of obstacles and targets). Identical workspace layouts were used for all trials in Session I so that subjects could enhance performance in subsequent trials by making modifications based on experience from previous trials. This was useful in training subjects on the remote cleanup task as well as in the training of subjects in how to design effective virtual fixtures.

Session II lasted from 10 to 15 minutes in which subjects were presented with a *new* workspace layout and were asked to perform a single trial of the remote cleanup task with no practice on this new task configuration. Subjects were required to take at least a 30 minute break between Session I and Session II so that they could approach the new task without lingering expectations from the initial testing. The objective of Session II was to explore whether subjects with minimal training in virtual fixture design, could design an effective fixture for a new task in an unfamiliar workspace. Whereas the initial workspace used in Session I represented a floor plan using only horizontal and vertical surfaces as obstacles, the new workspace used in Session II introduced rounded surfaces which would require a new approach to fixture design. This new layout, is shown below in Figure 7.9.

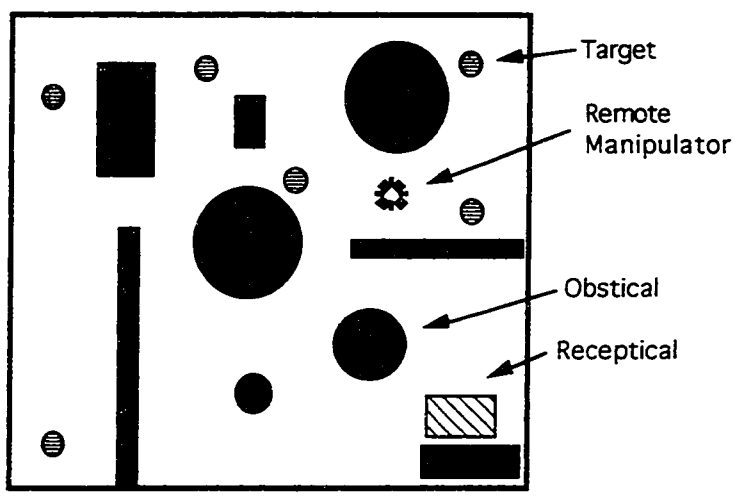

Figure 7.9. Remote Workspace for Cleanup Task II.

The two sessions were further divided into trial groupings called *tests*. Each of these tests were designed to introduce a new parameters into the experiment. Table 7.2 and Table 7.3 show the testing subdivisions for both Session I and Session II respectively. These tables enumerates the number of trials associated with each test. It should be noted that all tests were performed in the sequential order listed below except for Test 3 and Test 4 which were performed in the reverse order for half of the subjects. This was done to eliminate ordering effects in the presentation of virtual walls fixtures and virtual liquid fixtures so that their effectiveness could be compared.

Table 7.2: Testing Procedure for Session I

TEST #	DESCRIPTION OF TRIALS IN TEST	Trials
Test 1	**Practice Trials**: Subjects learn task with no virtual fixtures	6
Test 2	**Data Trial**: collect data on a trial with no virtual fixtures	1
Test 3	**Wall Trials:** Subjects introduced to virtual fixtures as Walls	3
Test 4	**Liquid Trials**: Subjects introduced to virtual fixtures as Liquids	3
Test 5	**Full Toolbox Trials**: Subjects introduced to full fixture Toolbox	3
Test 6	**Repeat Trial**: Subjects perform another trial with no virtual fixture	1

Table 7.3: Testing Procedure for Session II

TEST #	DESCRIPTION OF TRIALS IN TEST	Trials
Test 7	**New Task:** Subjects perform new task with fixture toolbox	1
Test 8	**New Task**: Subjects perform new task without virtual fixtures	1

7.2.4.2 No-Fixture Trials and Training Procedures

For both Session I and Session II, a single trial of the remote clean-up task consisted of a 90 second testing period in which subjects were instructed to try their best to generate the highest score they could. As part of the familiarization and training included in Session I, subjects were informed that a high score could only be generated by collecting targets quickly while at the same time accurately avoiding obstacles. Because this task required subjects to find their optimum balance between raw speed and controlled accuracy in hand motion, Session I began with a training period provided to all subjects so that they could fine tune their skill.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

The training session consisted of six practice trials to be executed without virtual fixturing. Fixtures were not introduced during initial training so that optimum performance in the basic task could be recorded for later comparison to fixture trials. This training period is lasted approximately 15 minutes. Subjects were shown their score at the end of each practice trial so that they could push themselves in subsequent trials in an attempt to maximize their performance. At the end of the six practice trials, the mean performance increase from one trial to the next across all subjects was under 2%. This small performance increase suggests that the subjects were very well trained and ready for data to be collected. Following this practice session, subjects were asked to perform a single trial of the cleanup task which would be recorded as their raw score for the experiment. Although this trial was used for data collection, it was run identically to the six preceding practice trials.

7.2.4.3 Testing Protocol and Session I Virtual Fixture Design

As listed in Table 7.2, Test sessions 3, 4, and 5 were used to introduce subjects to various fixture elements. Test 3 introduced subjects to virtual fixtures composed of only horizontal and vertical virtual walls. Test 4 introduced subjects to virtual fixtures composed of only viscous-liquid-regions. Test 5 introduced subjects to the full toolbox of virtual fixtures, including virtual walls, liquid regions, repulsive fields, haptic textures, and bumps. Before each of these tests, subjects were taught the simple keystrokes needed to define each virtual fixture element. Although subjects were taught the mechanism for creating fixtures, they were not provided with any strategies for designing effective virtual fixtures for the given task. This lack of influence in fixture design was done so that naive subjects could be observed in their fixture design process. It was hoped that by observing the natural inclinations of new users as they learn to use design effective virtual fixtures, insight would be gained about the process of using perceptual overlays to enhance manual performance. These observations were conducted by video-taping all fixture design sessions for later review.

During Tests 3,4 and 5, subjects were told that the first two trials were to be treated as practice during which time they were to design a fixture, perform a trial, and evaluate the effectiveness of their fixture design. At the completion of each practice trial, subjects were given a few minutes to look over their design and plan improvements for the subsequent trail. On the third trial of each test, subjects were told that it was no longer practice, it was the real task and that the results were being recorded as data. The two practice trials were provided to expose subjects to the notion of virtual fixturing and to familiarize them with the various perceptual elements in the toolbox before collecting actual data. Although subjects were not fully trained on fixture use after only two trials, data was collected because the intent of these studies was to demonstrate that users could design virtual fixtures with minimal training. Had more practice been provided, subjects would have eventually optimized their fixture design for the given task. Optimization

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

through continued practice however is not reflective of the intended application of the Perceptual Workstation.

Test 6 was the final test of Session I. It consisted of a single trial that was performed without the use of any virtual fixturing. This single no-fixture trial was used as a baseline to evaluate what part of the performance increases recorded in Tests 3,4, and 5 were due to learning rather than the introduction of virtual fixtures. By comparing this baseline trial to the initial no-fixture trial recorded in Test 2, the learning effect resulting from experience gained during Session I could be quantified. As it turned out, there was no measurable performance improvement in Test 6 over Test 2. This implies that no significant learning effect had occurred between these tests. This result suggests that the subjects were fully trained by the end of the practice period. It further suggests that performance increases recorded during the fixture trials could not be attributed to simple learning effects.

7.2.4.4 Testing Protocol and Session II

As shown in Table 7.3, Session II was composed of only two tests: Test 7 and Test 8. A minimum of a 30-minute rest period was required for all subjects between Session I and Session II in order to eliminate short term learning effects and reduce the impact of fatigue upon Session II results. Test 7 consisted of a single trial in which subjects were to design a virtual fixture using the full toolbox of virtual elements. This single trial was performed on a *new workspace* which had never been revealed to the subjects before. By requiring subjects design a virtual fixture in a new workspace, without practice or experience in that workspace, performance could be measured as an indication of how well users could design virtual fixtures from intuition and past experiences rather than from direct practice on the particular task at hand. Test 8 immediately followed Test 7 and consisted of a single no-fixture trial in the same new workspace. This no-fixture trial was the baseline against which Test 7 was to be compared. It should be noted that all subjects had the advantage in Test 8 of having actually performed in the new workspace once before during Test 7. Thus the performance increases revealed through baseline comparisons are likely to be slightly smaller than the true performance advantage. This mild ordering effect could not be avoided through trial randomization because the intent of the study was to have all subjects design fixtures for a workspace they had *never* performed in before. Thus, fixture trials had to precede the no-fixture baseline. Nevertheless, this study should provide a conservative indication of the performance advantage provided by virtual fixtures for tasks performed the very first time.

7.3. RESULTS

Six subjects were run through the full set of eight tests listed previously in Table 7.2 and Table 7.3. The results of these tests are tabularized below in Table 7.4 through Table 7.10. The scores listed were computed by tallying the number of targets successfully collected in the 90 second time period and subtracting the number of collisions with obstacles in the workspace.

Table 7.4: Test 1 Results: *Practice Trials without use of Virtual Fixtures*

Trial #	MEAN SCORE	PERCENT INCREASE
1	13.0	-----
2	17.1	32.1
3	20.1	17.5
4	19.1	5.0
5	21.3	11.3
6	21.5	0.8

Table 7.5: Test 2 Results: *Scores without use of Virtual Fixtures*

Trial #	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
1	17	20	23	26	20	24	21.6

Table 7.6: Test 3 Results: *Scores using the Virtual WALL Fixtures*

Trial #	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
1	28	36	36	38	45	32	35.8
2	32	38	43	46	41	36	39.3
3	40	39	44	47	44	36	41.6

Table 7.7: Test 4 Results: *Scores using the Virtual LIQUID region fixtures*

Trial #	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
1	18	22	30	32	33	22	26.1
2	23	24	30	28	32	24	26.8
3	24	25	31	30	30	30	28.3

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Table 7.8: Test 5 Results: *Scores using the Full TOOLBOX of Virtual Fixture elements*

Trial #	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
1	44	41	44	46	48	34	42.8
2	46	44	40	48	49	39	44.0
3	46	46	44	46	50	38	45.0

Table 7.9: Test 6 Results: *Baseline Scores without the use of Virtual Fixtures*

Trial #	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
1	22	19	22	27	20	19	21.5

Table 7.10: Test 7 & 8: *Scores in New Task With and Without Virtual Fixtures*

	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	MEAN
With Virtual Fixtures	31	30	34	35	42	27	33.2
Without Virtual Fixtures	18	18	24	23	25	14	20.3

7.4. ANALYSIS

To quantify the effect that virtual fixturing had upon operator performance, the *percentage performance improvement* was computed for all trials. In both Session I and Session II, percentage performance improvement was computed as the percent increase in virtual fixture trial scores as compared to the no-fixture baseline scores. Session I results compared fixture scores to the Test 2 baseline. Session II results compared fixture scores to the Test 8 baseline. The results of these analyses are listed below in Table 7.11.

Table 7.11: Performance Improvement of Fixture Trials over Non-Fixture Trials

TEST: FIXTURE USED IN TEST	Mean Performance Improvement
Test 3: Virtual Wall Fixtures	96 % ± 28%
Test 4: Virtual Liquid Regions	33 % ± 12 %
Test 5: Full Toolbox in known Workspace	109 % ± 40 %
Test 7: Full Toolbox in unknown Workspace	66 % ± 18%

7.5 DISCUSSION

While previous experiments have confirmed the notion of virtual fixturing by testing performance in simple standardized peg-insertion tasks, this set of experiments was the first exploration of the virtual fixture concept in which a complex and challenging remote task was implemented. Although convenient for quantifying human performance in a controlled testing scenario, simple peg-insertion trials do not represent particularly appropriate applications of the virtual fixturing concept. The present set of experiments implemented a remote hazardous cleanup paradigm which provided a significantly more complex and challenging telemanipulation scenario. The results of these hazardous cleanup trials confirm that virtual fixtures were highly effective in enhancing operator performance. As is listed in Table 7.11, the results of Tests 3, 4, 5, and 7 have revealed significant performance improvements ranging from 33% to 109%. The fact that virtual fixtures were used effectively in this complex telemanipulation scenario suggests that the results of the previous experiments, which implemented simple peg-insertion trials, can now be generalized to more complex and realistic paradigms.

All previous testing of the virtual fixture concept had implemented a standard set of fixtures designs whose configurations were pre-defined by the experimenter. Thus, a primary question to be answered by the current set of experiments was as follows: Can naive subjects design and implement their own virtual fixtures and enhance their own performance in telemanipulation tasks? As listed in Table 7.11, the results of these experiments strongly confirm that subjects were able to design very effective virtual fixtures to enhance their personal performance in the given remote cleanup task. In fact, the results of Test 5 revealed that when confronted with a telemanipulation task which they were familiar with, test subjects were on average able to *more than double* their performance by implementing virtual fixtures they had designed themselves using the full toolbox of fixtures elements. When confronted with a new workspace layout (one in which they had never encountered before), subjects were still able to design virtual fixtures which enhanced their personal performance by an average of 66%. These dramatic performance improvements suggests not only that the naive subjects were comfortable using the virtual fixture metaphor to develop perceptual overlays, it suggests that subjects quickly became skilled at crafting very effective fixture designs. The 66% performance improvement recorded for trials performed in an unknown workspace suggests that subjects quickly developed an *intuition* about fixture design which allowed them to construct effective perceptual overlays for tasks which they had never actually performed. This result implies that virtual fixtures can be useful in telemanipulation environments whose layouts and configurations are not well known in advance.

The fact that all subjects were able to design effective virtual fixtures for a new workspace which they had never actually experienced, suggests the existence of some generalized rules or guidelines to fixture design which can be extracted and applied generically to a given task in a given workspace. Although all subjects apparently acquired some generic guidelines for fixture design, the only formal instruction subjects received was on the mechanical *procedure* by which virtual fixtures are created. Subjects were not given any instruction or guidelines on the strategy of fixture design. Thus all subjects learned to design effective virtual fixtures through the personal experience accrued during a small number of practice trials. Although subjects performed Test 5 with only the minimal experience of 8 previous fixture trials, the mean performance improvement was 109%. This is a very significant performance improvement for subjects with so little training and experience designing virtual fixturing. It is likely that with substantially more experience with fixture design and usage, the performance improvement resulting from fixture use would be even greater.

To gain insight into the generic strategies for fixture design that subjects developed during the testing sessions, all fixture design sessions were video taped for later observation. Subjects were also given post testing interviews to provide further insight into the design and implementation of various fixture configurations. An informal review of the video tapes and post testing interviews has revealed that all subjects eventually adopted one of two basic strategies for fixture design. One strategy was to design fixtures which used rigid walls, repulsive fields, haptic textures, and/or viscous-liquid-regions to define *barriers* which prevented collisions with the obstacles. Post testing interviews revealed that subjects who designed fixtures using this strategy felt more confident moving safely through the workspace when barriers were in place. This confidence was reported even when the barriers defined by fixturing components did not provide actual kinematic protection against contacting the obstacles. Thus we can look upon fixture barriers as providing both perceptual and physical protection from obstacle collisions. Subjects reported having to devote less visual attention to avoiding obstacles when barriers were in place. As a result, subjects were able to devote more visual attention to planning the paths to-and-from targets. Even when the barriers generated by the virtual fixturing could be penetrated, as is true of viscous-liquid-regions and haptic textures, subjects reported being able ease the visual attention devoted to obstacle avoidance because the additional haptic cues alerted them to imminent collisions.

Although rigid walls, viscous-liquid-regions, and haptic textures, were all effectively used as barriers to prevent collision with obstacles, rigid walls were greatly preferred by all subjects tested. Most subjects experimented using multiple fixture elements as barriers during the practice trials, but as they progressed in experience they eventually abandoned all but the rigid walls. An example of a virtual fixture composed by the strategy of building barriers around the obstacles is shown below in Figure 7.10. As is implied by the figure, a virtual fixture designed using this strategy results in a perceptual overlay which is

useful in aiding subjects in the avoidance of obstacles, but provided subjects with little direct assistance in acquiring targets or planning trajectories.

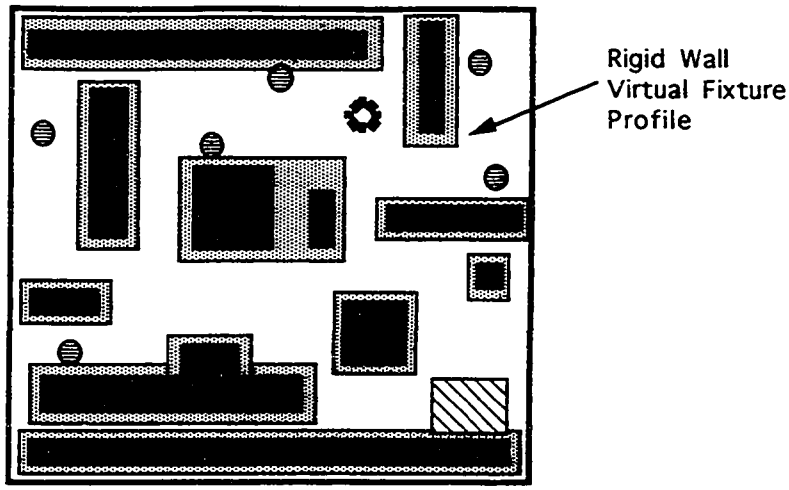

Figure 7.10: Remote Environment with Overlaid Virtual Fixture as Obstacle Barrier

The other basic strategy for fixture design which many subjects adopted was to build virtual fixtures which defined a path through the workspace, rather than outlining the obstacles to be avoided. Such a strategy was used to define a trajectory, free of obstacles, which could be traversed very rapidly using haptic cues to guide the joystick. All subjects who adopted this strategy, also used the fixture as a means of target acquisition by designing the path such that it would guarantee contact between the remote manipulator and the target in question. Although subjects used multiple fixture elements in such fixture designs, the rigid wall was by far the preferred element for defining the boundaries of the path. All subjects who used this strategy would outline the desired path using the rigid wall contacts, and then embellish the path with additional perceptual information in the form of repulsive fields, bumps, textures, or viscous liquid regions. For example, two subjects designed virtual fixtures by first defining a path using rigid walls and then by placing small bumps in front of the targets. Post testing interviews revealed that such bumps were placed in front of the targets so that the subjects could "feel the bump as a signal" which indicated contact with the target. Thus, the addition of the haptic bumps to the fixture design allowed the task to be completed with less visual attention being devoted to target acquisition. This act overlaying abstract haptic cues on top of the workspace in order to enhance performance and reduce the dependence on taxed sensory modalities, demonstrates that the fixture metaphor allowed subjects to fully embrace the concept of perceptual overlays in its intended form.

An example of a virtual fixture which was designed by this strategy is shown below in Figure 7.11. It was composed by first defining a path for the remote manipulator and was then embellished with additional haptic cues. Included in this fixture design are rigid walls which define the path and bumps which signal that a target has been reached. Also included in this design are repulsive fields which are positioned to influence the operator's trajectory by biasing motions away from sharp corners. This repulsion from sharp corners was used by a number of subjects to achieve smooth manual operations in an environment in which curved walls could not be defined.

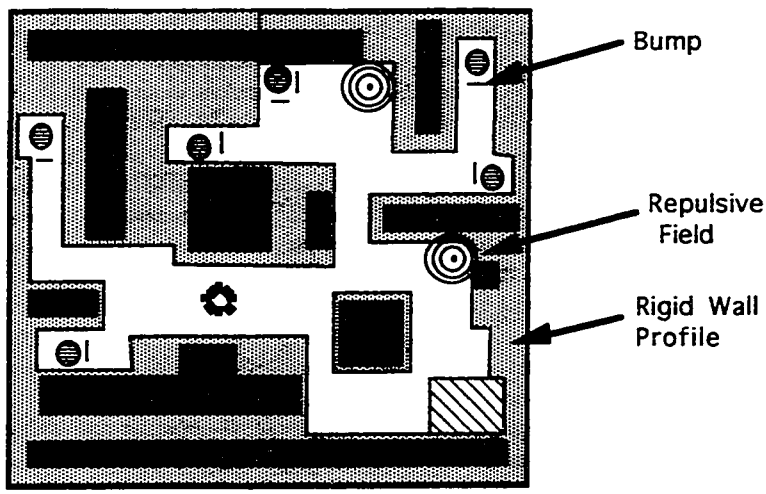

Figure 7.11: Remote Workspace with Overlaid Virtual Fixture to define Path to Targets

Another important goal of this set of experiments was to answer the basic question: Can test subjects make use of a *toolbox* of perceptual elements to compose virtual fixtures? The performance improvements of 109% and 66% during Test 5 and Test 7 confirms that the toolbox of perceptual elements was effectively used. This result confirms that an interactive design environment, in which virtual fixtures are built-up from basic elements chosen from a toolbox, is a viable methodology for allowing users to create their own perceptual overlays for a given task. This serves as a powerful confirmation of the perceptual workstation concept and suggests that the Percept-Station system is a valuable initial model for such a workstation.

7.6. CONCLUSIONS

This study has extended the results of previous experiments and strengthened the overall support for the virtual fixture concept. While previous studies demonstrated the effectiveness of virtual fixtures on a simple set of controlled peg-insertion trials, this study was the first implementation of virtual fixturing in a complex and challenging manipulation paradigm. Results of this study have shown that virtual fixturing was highly effective in enhancing manual performance in a remote hazardous cleanup task.

All previous tests of the virtual fixture concept have implemented a pre-defined set of virtual fixtures which were designed and constructed by the experimenter. The present set of tests was the first study of virtual fixturing in which subjects were required to design there own virtual fixtures to enhance their performance in a given task. The resulting performance increases, which were as high as 109%, demonstrated that naive subjects were highly capable of designing effective virtual fixtures for the given task. This result supports the concept of a Perceptual Workstation in which users can interactively design and implement virtual fixtures when confronted with new or changing scenarios. These results also confirm that providing users with a toolbox of perceptual elements is an effective means of allowing subjects to design virtual fixtures from basic perceptual building blocks.

The final trials in this study required subjects to design virtual fixtures to enhance their performance in a remote workspace which they had never experienced. The resulting 66% mean performance increase across all subjects strongly suggests that with only minimal training, subjects developed general guidelines for fixture design which could be effectively applied to unknown environments. This result supports the potential use of a Perceptual Workstation in which subjects can interactively design virtual fixtures when confronted with new or changing tasks.

An informal observation of fixture design methodology has revealed that all subjects adopted one of two basic strategies when designing virtual fixtures. One strategy was to define barriers around obstacles in the workspace. It is believed that this strategy was effective because it allowed subjects to maneuver through the remote environment while devoting less visual concentration to collision avoidance and more visual attention to path planning and target acquisition. The second strategy was to use virtual fixture elements to actually define the path between targets and the receptacle. It is believed that this strategy was effective because it provided overlaid perceptual information which both assisted subjects in the avoidance of obstacles but also helped in the planning of paths and the acquisition of targets. Such path-like virtual fixtures were often embellished with sensations such as bumps and repulsive fields which enhanced performance by introducing additional haptic cues.

Chapter Eight

"A major role of new technology should be to make tasks simpler... "

- Donald A. Norman
The Design of Everyday Things

---

8. Summary and Beyond...

This thesis addressed a wide range of topics including virtual reality, telepresence, human perception, human performance, human-computer interaction, and design. The notion of virtual reality was explored from the perspective of a designer trying to implement this new technology in a useful application. The concept of *Virtual Fixtures* was introduced in this thesis and empirically tested under a wide variety of conditions. Extensive subject testing demonstrated the effectiveness of the fixturing concept a well as provided insights into why such perceptual overlays enhance human performance. In the process of developing and testing the potential of virtual fixturing, this thesis borrowed many ideas from the field of perceptual psychology and attempted to formulate a theoretical foundation for the design of virtual percepts. A preliminary design methodology called *Design-for-Perception* was presented and explored. Empirical tests of the virtual fixture concepts provided experimental insights into the principles of Design-for-Perception. Although Design-for-Perception was addressed within the context of virtual fixturing, its implications extend to the design of general immersive human interface systems. In addition to presenting theoretical design ideas, this thesis introduced and tested a practical design procedure called *Perceptual Decomposition* which was intended to aid in the perceptual designer in the construction of virtual percepts. Finally, a limited prototype of a *Perceptual Workstation* was developed to demonstrate that the theoretical ideas behind virtual fixturing and Design-for-Perception can be extended to a real-world application. Empirical testing demonstrated that human operators could design their own virtual fixtures from a toolbox of perceptual elements when confronted with a given task. Such fixtures were found to be highly effective in enhancing operator performance in a complex manual operation.

The following sections give a brief summary of the major results and conclusions revealed throughout this dissertation. Recommendations for future development and experimentation are also presented here. The intent of this chapter is not to reiterate all of the results listed in preceding chapters. Rather, the following sections are intended to give a concise summary of the major points.

8.1 Design-for-Perception

As introduced in Chapter 1 and explored throughout the rest of this thesis, Design-for-Perception is methodology for the design and implementation of immersive human interface systems such as virtual reality and telepresence. In many ways, Design-for-Perception is more of a *mind-set* for design than it is a formal procedure for design. The fundamental principle is that virtual percepts and immersive human interface systems be designed from a *perceptual* rather than *physical* perspective. It asks designers to take a step back and consider philosophical and psychophysical issues relating to the notion of *reality* before setting out to design a *virtual reality*. Design-for-Perception suggests that designers maintain a strict separation in their thinking between *physical reality* and *perceptual reality*. Such a mind-set will help the designer overcome the natural human tendency to equate *reality* with the *physical world*. Known as *naive realism* by perceptual psychologists, the tendency to assume that our perception of the world is a direct representation of our physical surroundings can lead to highly inefficient presentations of virtual information. Design-for-Perception is particularly applicable to the generation of virtual *haptic percepts* because limitations of force reflecting hardware make it impossible to produce accurate physical models of most haptic phenomenon.

When developing the virtual representation of a sensory percept, such as the haptic sensation produced when encountering a rigid wall, the methodology of *Design-for-Perception* suggests little need for producing an accurate physical model of the *distal stimulus*. Rather, *design-for-perception* prescribes generating a perceptually adequate model of *proximal stimulus*. When such a model of the proximal stimulus is provided, the user should have little trouble drawing the correct *perceptual hypothesis* and inferring the appropriate distal stimulus. Although such a procedure is likely to produce an efficient representation of a virtual percept, it does pose some unique design challenges. The primary challenge is that in order to produce a *perceptually adequate* representation of a virtual percept, a strong understanding of the perceptual qualities of the percept being modeled is required. Such a perceptual understanding might be acquired through experience and training as is true of the visual artist learning to *sketch* the salient perceptual features of a visual scene. Unfortunately, when it comes to designing haptic percepts, there are

no formal "sketching techniques" to help the perceptual designer generate effective and efficient representations.

The approach described in this thesis for exploring the perceptual content of a haptic sensation involved *decomposing* the given percept into its salient perceptual features. An initial attempt was made to perform the "perceptual decomposition" of a *rigid-surface* sensation through observation and introspection. The author spent many hours interacting with a two degree of freedom (dof) force reflecting joystick, gaining insight into the feel of a wide variety of virtual haptic representations. The goal of such an exploration was to isolate distinct perceptual qualities of a rigid surface and reveal how to reproduce such perceptual qualities through simple virtual models. Although perceptual modeling based on personal observation was insightful, it was the subjective result of exploration rather than objective experimentation. Thus a formal procedure was developed which involved the systematic testing of naive test subjects. As described in detail in Chapter 4, this process of *Perceptual Decomposition* recorded the subjective impressions of test subjects interacting with various virtual haptic sensations. The results of such testing confirmed that a rigid-wall percept could be decomposed into discrete perceptual qualities. It further confirmed that effective virtual representations of haptic percepts could be constructed by presenting an accurate perceptual model of the salient perceptual features.

In an interesting confirmation of the Design-for-Perception methodology, the most effective virtual representations constructed through perceptual decomposition were very different from what physical modeling would have predicted. Virtual rigid-surfaces derived through *physical modeling* are generally composed of a very high stiffness in combination with a small amount of added damping to provide stability. The most effective rigid-surface derived in Chapter 4 through *perceptual modeling* was composed of a very high damping to provide the illusion of an abrupt initial contact (a sharp jolt) plus just enough stiffness to maintain an illusion of rigidity. The perceptual importance of a crisp jolt (as opposed to static hardness) is likely a result of the fact that the human sensory system is generally more sensitive to differential changes than to static values. Thus the perceptual representation of a rigid-surface is not only more effective than a physically derived representation, but its content could not have been produced through traditional physical modeling techniques which do not account for human perceptual requirements.

8.2 Virtual Fixtures

A primary outgrowth of the Design-for-Perception methodology was the idea of enhancing human performance by embellishing the user's *perceptual reality* with abstract sensory information known as *Perceptual Overlays*. Because the abstract notion of perceptual overlays is as difficult to conceptualize as it

is to discuss, the *virtual fixture metaphor* was introduced to describe these overlays as concrete entities. Virtual Fixtures were defined as follows:

> **Virtual Fixtures**: Abstract perceptual information overlaid on top of the reflected sensory feedback from a remote environment such that a natural and predictable relation exists between an operator's kinesthetic activities (efference) and the subsequent changes in the sensations presented (afference).

Thus virtual fixtures are computer generated percepts overlaid on top of the sensory feedback from a remote telepresence worksite. The operator of a telepresence system will feel present with respect to both the *real* remote environment and the *virtual* overlays. Because virtual fixtures are independent of all feedback from the remote environment, they are immune from the communication delays and bandwidth limitations which distort and disrupt the perceptual information reflected from the remote site. Thus in many ways, virtual fixtures are actually the highest fidelity representations within the operator's perceptual reality.

Like the ruler guiding the pencil, virtual fixtures overlaid on top of a remote workspace can reduce mental processing required to perform the task, reduce the work load of certain sensory modalities, and most of all allow precision and performance to exceed natural human abilities. Although virtual fixtures could be functionally equivalent to fixtures in the real world, there are many advantages inherent to virtual fixtures because they are computer simulations as opposed to real physical hardware. When overlaid on a workspace, the fixtures only interact with the user and not with the workspace itself. Thus fixtures can occupy the same physical space as objects in the remote environment. This means that the workspace geometry imposes no constraints upon the placement or configuration of virtual fixtures. In addition, such fixturing has no mass, has no physical or mechanical constraints, requires no machining time or maintenance, can be easily prototyped and modified, and can essentially be transported to remote locations using nothing more than standard communication links.

8.2.1 Experimental Results

Through a series of empirical tests, virtual fixtures were found to enhance human performance in standard manual tasks. In Chapter 2, virtual fixtures were tested as ideal virtual representations, constructed out of actual physical hardware rather than being computer generated. Such testing revealed up to 57% mean performance increase across all subjects when using a particular haptic fixture design. Chapter 2 also tested the *cross-modal* combination of haptic and auditory fixtures and found a mean performance increase across all subjects of up to 70%. Analysis revealed that the addition of simple auditory perceptual overlays such as sonic fields, enhanced operator performance by as much as 23% over haptic fixtures alone.

In Chapter 3, virtual fixtures were investigated in the context of time-delayed telemanipulation. It was found through human testing that time-delays as small as 450ms could degrade operator performance by

as much as 44%. Using particular virtual fixture designs, it was found that the performance degradation due to time-delay could be completely eliminated. In fact, mean movement-times (across all subjects) for peg-insertion trials using certain virtual fixtures at 450ms time-delay were actually faster than no-fixture trials performed with no-delay. Since time-delay is a primary technical challenge of telemanipulation systems, the effectiveness of virtual fixtures to function under time-delay conditions is a strong affirmation of the fixturing concept.

In Chapter 5, virtual fixtures were presented as true computer-generated constructs using a six degree of freedom upper-body exoskeleton. Despite bandwidth limitations of the force reflection hardware, virtual fixtures were generated which enhanced operator performance by as much as 102%. These tests were the first application of abstract haptic fixture designs in three-dimensional space. These experiments served as a basic confirmation that the virtual fixturing concept remains effective under non-ideal conditions. These experiments also extended the confirmation of the fixturing concept to three dimensions. In Chapter 6, audio virtual fixtures were presented in three-dimensional space as audio perceptual overlays composed of 3D binaural localization cues. These abstract fixture designs were presented through an Audio Localization Cue Synthesizer. Subject testing confirmed that binaural perceptual overlays could enhance operator performance by as much as 26%.

8.2.2 Insights and Implications

In addition to confirming the effectiveness of virtual fixtures under a variety of testing conditions, the extensive human testing described throughout this thesis also provided insights into *why* virtual fixtures enhance human performance. Human performance data, post testing interviews with subjects, and recorded task trajectories were all considered when evaluating the reasons behind fixture effectiveness. The following paragraphs summarize the most promising conclusions drawn.

Virtual Fixtures introduce abstract perceptual information into a user's perceptual reality. It has been hypothesized that such overlaid percepts provide users with additional means through which to relate afference to efference. Current theories in perceptual psychology suggest that lawful relations between afference and efference contribute greatly to the overall illusion of *presence* within an immersive environment. Thus the perceptual overlays introduced by virtual fixtures may foster operator presence within telepresence systems. In addition, the introduced sensory information provided by virtual fixtures may allow neural processing through multiple sensory modalities. As a result, the added sensory information provided by virtual fixtures may act to increase an operator's information processing capacity (through-put) and therefor increase telepresence performance. Such overlaid percepts can also be effective in displacing the processing burden from taxed sensory modalities. For example, a standard peg-insertion task is visually intensive. With the use of haptic or audio virtual fixtures, visual inspection can be replaced by

haptic cues or auditory signals. This might free up the visual modality for other activities or might simply allow the processing burden to be shared in a more efficient way. Overlaid perceptual information may also provide users with additional means by which to gauge task performance. For example, a standard peg insertion task might simply be monitored as "looking a certain way". With the introduction of overlaid haptic and audio percepts, the task might also be gauged as "feeling a certain way" and "sounding a certain way". Such additional criteria by which to optimize task execution is likely to enhance operator performance in telepresence tasks. Finally, in time-delayed teleoperation, virtual fixtures were found to provide real-time perceptual cues into an otherwise time distorted perceptual reality. Such real-time reference points allowed users to abandon typical *move-and-wait* strategies during telemanipulation with time-delay.

In addition to the perceptual effects described above, virtual fixtures may also enhance performance by physically influencing operator activity. For example, haptic percepts such as rigid surfaces may provide physical constraints which reduce the degrees-of-freedom of a user's perceptual reality. This could enhance performance by decreasing mental workload because fewer alternatives need be considered during task execution. Such physical constraints may also reduce the number of possible failure modes or eliminate inefficient alternatives from the task-space. Finally, virtual fixtures may be used as protective barriers to prevent against dangerous or destructive actions. This could enhance performance by reducing the criteria for success or by simply boosting an operator's confidence.

8.2.3 Virtual Fixtures and Theories about Task Difficulty

In his landmark book, *The Design of Everyday Things* (also published as *The Psychology of Everyday Things*), Donald Norman applies the principles of cognitive psychology to the design of effective products. As a guideline for enhancing human performance through effective design, he proposes "seven principles for transforming difficult tasks into simple ones". Although *virtual fixtures* are very different from the *consumer products* envisioned by Donald Norman, they are entities which are intended to enhance human performance in manual tasks. In the following subsections, the reasons for fixture effectiveness listed above will be correlated with the seven design guidelines proposed by Norman. This is a constructive exercise because it presents the potential of perceptual overlays in terms of well known design principles.

8.2.3.1 "1. Use both knowledge in the world and knowledge in the head"

Donald Norman suggests that people perform better when the knowledge required for a task is "explicit in the world or readily derived through constraints". Virtual fixturing is an excellent means of

making the knowledge required for telepresence tasks explicit in an operator's perceptual reality, even through such knowledge may not be explicit in the remote worksites. For example, a perceptual overlay composed of haptic surfaces could define the proper execution of a task, making it explicit in the user's perceptual reality. Norman suggests that performance will improve once the knowledge in the world, becomes internalized in the user's head. He therefor suggests that designers present the explicit knowledge about task execution in a "natural format" which is easily internalized. The virtual fixture metaphor complies well with this design guideline by presenting information in so natural a format, that fixtures are actually perceived as genuine physical components of the task environment.

8.2.3.2 "2. Simplify the structure of tasks."

Norman says that tasks should be presented as simply as possible, minimizing the amount of planning or problem solving they require. He says that technological innovations should be used to restructure complex tasks into simpler ones. This is particularly true for tasks requiring high dexterity or manual skill. Virtual Fixtures, by their very nature, simplify the structure of manual tasks. For example, one reason why a peg-insertion task is challenging is because the user needs to maintain an efficient trajectory. This means not retracting the peg too far from the task-board because that wastes time. This also means not getting too close to the task-board because a collision is an even bigger waste of time. Thus maintaining proper distance from the board is a complex and challenging manual operation. By introducing perceptual cues which explicitly signals if the user is too close or too far from the task board, the task requirements become simpler and performance is improved.

8.2.3.3 "3. Get the mappings right"

To enhance human performance, Norman suggests that tasks be represented through "natural mappings". By this he means that the relationship between actions and outcomes should be as direct as possible. He goes on to say that feedback must be provided which is in a format that is easy for the user to understand and which directly relates to the user's intentions. In light of such requirements, perceptual overlays in the form of virtual fixtures provide an ideal medium for user feedback because they convey information in a very natural and intuitive format. For example, when a user moves his or her hand into a haptic fixture composed of a virtual wall, the resulting jolt may be abrupt, but it is completely expected. The implications of the jolt, as described by its magnitude and duration, are easily interpreted by the operator whom, like all people, is well versed in interacting with rigid walls.

8.2.3.4 "4. Make things visible."

According to Norman, performance can be improved by making things *visible* to the user which are not ordinarily so. He uses the speedometer in an airplane as an example, making the velocity of the

vehicle visible to operator. Providing such information to a user can facilitate the generation of an effective mental model of the given task. Unfortunately, the use of dials and gauges are not a particularly natural or intuitive means of presenting information. Perceptual overlays, on the other hand, make inherent properties of task execution "visible", without resorting to un-natural modes of presentation such dials and gauges. Perceptual overlays present the information through a natural means that can be rapidly internalized. For example, in Chapter 2, auditory fields were presented to test subjects. While using such virtual fixtures, subjects reported having a better sense of hand dynamics because motion no longer just *looked* and *felt* a certain way, but hand motion also *sounded* a certain way.

8.2.3.5 "5. Exploit the power of constraints, both natural and artificial"

In order to enhance task performance, Norman suggests that designers implement constraints that reduce the number of alternatives which the operator has to consider. Virtual Fixtures composed of haptic overlays do just that. Consider the fixture designs implemented in Chapter 5 which were composed of two intersecting haptic surfaces. One surface was a *back-plane* which kept the operator from pulling too far from the task-board. The other was a *target-stop* which kept the user from overshooting the target-hole. Both of these elements constrained the user from performing an inefficient trajectory subsequently enhanced performance.

8.2.3.6 "6. Design for Error"

Designer should expect users to make errors in task execution. Therefor, Norman suggests that effective designs should account for such errors and allow users to recover easily. He suggests using *forcing-functions* rather than linear pathways of progressive steps because they are easier to recover from if an error occurs midway through a task. All actions should be reversible. These too are properties inherent to the design of virtual fixtures. Even highly constrained fixtures such as the snap-to-lines and snap-to-planes introduced in Chapter 5 are flexible enough to support user error. Users maintain the ability to "snap-in" or "snap-out" of the fixture at will.

8.2.3.7 "7. When all else fails, standardize"

The use of standard designs enhances operator performance for the simple reason that things only need to be learned once. To take advantage of this principle in the design of virtual fixtures, it has been proposed that a *toolbox* of standard perceptual elements be provided to users when designing virtual fixtures. Although this thesis does not claim to define which standard elements should be included in such a toolbox, testing has shown the toolbox concept to be a viable method of providing standard perceptual building blocks for fixture design.

8.3 Perceptual Workstation

Although virtual fixturing is a powerful concept, if the development of effective fixtures requires complex computation or intimate prior knowledge of the remote workspace, robot autonomy might be a preferred solution to the task at hand. If, on the other hand, effective fixtures could be rapidly developed by the human operator of a telepresence system through the use of an interactive fixture design interface, virtual fixtures could be used in unstructured or changing environments unsuitable for autonomous systems. If we consider that the basic goal of telepresence systems is to extend an operator's natural manual and cognitive abilities to a remote environment, it seems that providing the ability to interactively design and modify assistive virtual fixtures is essential to preserving the flexibility of the human in the loop.

As described in detail in Chapter 7, a limited prototype *Perceptual Workstation* was developed to allow operators of telepresence systems to design and implement assistive virtual fixtures when confronted with a given task. The prototype system provided users with a *toolbox* of basic fixture elements from which effective virtual fixtures could be quickly designed and implemented. The toolbox allowed subjects to design fixtures composed of haptic percepts such as rigid walls, repulsive fields, viscous liquid regions, texture gratings, and bumps. Empirical testing of naive human subjects has confirmed that operators can design and implement their own virtual fixtures when confronted with a task. Using perceptual overlays that they designed themselves, the subjects in these experimental sessions were able to more than double their performance in a given task.

8.4 Recommendations for Future Development and Experimentation

Because this thesis covers a wide range of topics and introduces many new concepts, the questions raised by these pages far outnumber the questions answered. Thus, I will not even attempt to use this section to address all of the issues which I feel deserve future study. I will however, enumerate what I consider to be the few most important *next-steps* in the exploration of the virtual fixturing concept.

What may be the most neglected facet of virtual fixturing thus far explored is the use of the *visual modality* as a medium for virtual fixture generation. The studies presented in this thesis did not take advantage of the visual modality because of the technical challenges associated with overlaying computer generated images on top of images reflected from the remote environment. Although possible with off the shelf hardware, the equipment cost for overlaying 3D graphics on top of 3D video were prohibitive for the

studies performed as part of this dissertation. Nevertheless, the use of monocular and/or stereoscopic visual cues offer a wide range of fixturing possibilities and should be pursued with high priority.

What may be the most interesting phenomenon yet to be explored in the context of virtual fixturing is the effect of *cross-modal combinations* of sensory percepts upon human performance. Most studies performed in this thesis only implemented perceptual overlays targeting a single sensory modality. One study (as described in Chapter 2), did implement cross-modal combinations of haptic and audio percepts and the results far exceeded the use of haptic or audio sensations used in isolation. Intuition suggests that the true power of virtual fixturing will not be revealed until such compound virtual fixture designs are fully explored. When cross-modal fixtures are investigated, other perceptual design issues will be raised such as the required accuracy of registration between sensory percepts. For example, how well do the haptic and visual representations of a virtual rigid wall need to be in registration for an effective percept to be achieved. Design-for-Perception principles suggest that such registration requirements have little to do with physical accuracy and much to do with human perceptual requirements. Thus far, a single cross-modal registration study (haptic-audio registration in a virtual rigid wall) has been explored, but the results of this study were excluded from the main text of this thesis for brevity. A reprint of a short conference paper on this work is included in the Appendix for reference. Although this work constitutes a first step towards addressing perceptual requirements of cross-modal virtual percepts, an entire series of studies should be performed in order to fully explore cross-modal combinations of haptic, visual, auditory, and even vestibular sensations.

Similar to the perceptual requirements for cross-modal registration, the perceptual requirements for virtual fixture calibration is an issue which deserves future study. Although the spatial and temporal registration requirements for perceptual overlays are different for every task and every fixture design, studies should be performed to reveal how well a virtual overlay needs to be calibrated to a real environment. Although no formal studies have been performed, experience has revealed that very accurate spatial registration was *not* required for effective implementation of virtual fixtures. However, it does seem logical that large calibration errors will hinder certain types of virtual fixture designs such as the Snap-to-Line fixtures presented in Chapter 5. On the other hand, fixture designs such as the Uniform Damping Field presented in Chapter 5, have essentially zero requirements for spatial calibration. Nevertheless, some means of quantifying the required registration accuracy for various fixture designs would be a valuable exorcise.

Additional experiments should also be performed to expand the *perceptual decomposition* studies described in Chapter 4. Although the initial perceptual analysis of the rigid surface percept was an insightful and instructive exercise, continued testing should be done to improve the virtual representation of the rigid surface percept as well as to optimize the process of perceptual decomposition itself. For example,

initial studies revealed an interesting perceptual relationship between stiffness-to-damping ratio and the effectiveness of a virtual surface model. Future studies should investigate a wide range of stiffness-to-damping ratios in order to fully explore this effect.

Once a firm perceptual understanding of the simple rigid surface is attained, deeper investigations into the perceptual representations of more complex virtual percepts will raise many additional research questions. Just taking a small step beyond the simple rigid surface, issues such as surface material, surface texture, even surface thermal properties are raised. This brings up many interesting questions which could be answered though Design-for-Perception. For example, how does one create a virtual representation of a *steel-surface* and how is it different from a virtual representation of *glass-surface* or a *wood-surface*. Rather than looking at the physical properties which differentiate these surfaces, we could decompose the perceptual differences. Studies could be done to reveal the primary perceptual cues which imply steel-ness or glass-ness or wood-ness. Such applications of the Design-for-Perception methodology would hopefully result in efficient and effective virtual representations as well as lay the foundation for a design vocabulary for virtual entities.

One final area which deserves future study is the effectiveness of virtual fixtures not only to enhance net performance, but also to enhance the rate of *training* in manual tasks. Simple studies could be performed which quantified the rate of learning for manual tasks performed with and without the use of perceptual overlays. Although the mechanics of such studies would be simple, variability in learning rates between users would likely require a very large subject pool. For this reason, learning based experimentation would be an extensive undertaking. Nevertheless, there is much reason to expect virtual fixtures to flatten the learning curve on manual tasks. It is well known that *feedback* has a major impact on human learning, and perceptual overlays act to increase the modes of feedback accessible to the user. Also, virtual fixtures can be designed which impose constraints on manual task. Such constraints could be chosen to reduce the modes of failure or reduce alternatives for success. This type of task simplification would be expected to increase the rate of learning. There is also the possibility of developing *sequences* of fixture designs which enhance operator performance through in progressive steps. Such virtual "training-wheels" could evolve over time into a highly efficient virtual fixture or could be ultimately abandoned once the task was adequately learned. Learning effects also imply that additional studies be performed to investigate the *residual effects* of fixture use. Tests should be performed to quantify whether subjects whom undergo extensive practice using virtual fixtures retain any of the performance improvement when fixtures are removed.

Clearly the notion of virtual fixturing and the methodology of design-for-perception have opened the doors to a wide range of future studies. In addition, there are a vast array of potential applications for

virtual fixturing from tele-surgery to immersive computer-aided-design which should be explored. I hope the previous eight chapters act as a call to arms, motivating designers to employ perceptual tools such as virtual fixturing to enhance human performance and understanding within immersive human interface systems. I also hope that the previous eight chapters draw attention to the importance of perceptual parameters upon the design of virtual percepts and encourage designers of virtual reality systems to avoid slipping into naive realism.

BIBLIOGRAPHY

Adelstein, B.D., *A Virtual Environment System for the Study of Human Arm Tremor.* Ph.D. dissertation, Dept. of Mech. Eng., M.I.T., Cambridge MA, 1989.

Adelstein, B.D., and Rosen, M.J. A high performance two degree-of-freedom kinesthetic interface. "Human-Machine Interfaces for Teleoperators and Virtual Environments", NASA Conference Publication 10071, pp. 108-113, 1991.

Adelstein, B.D., and Rosen, M.J. "Design and implementaion of a force reflecting manipulandum for manual control research." In *Advances in Robotics*, ed. H. Kazerooni, Amer. Soc. Mech. Eng, New York, pp. 1-12, 1992.

Akin, D.L., Minsky, M.L., Thiel, E.D. and Kurtzman, C.R. (1983). "Space Applications of Automation, Robotics, and Machine Intelligence Systems," (ARAMIS)- Phase II, Vol. 3: Executive Summery, M.I.T., Cambridge, MA.

Arnheim, Rudolf. *Art and Visual Perception*, University of California Press, Los Angeles, California. 1974.

Audio Localization Cue Synthesizer (ALCS) User's Manual, Armstrong Laboratory, Wright-Patterson AFB, OH. March 1992

Bartholet, Stephen J., "An exoskeleton Master Arm, Wrist, and End Effector Controller with Force Reflecting Telepresence," U.S. Department of Defense SBIR Phase I Final Report, Contract No. F33615-89-C-0587, May 1 1989.

Bartholet, Stephen J., "An exoskeleton Master Arm, Wrist, and End Effector Controller with Force Reflecting Telepresence," U.S. Department of Defense SBIR Phase II Final Report, Contract No. F33615-89-C-0587, 1991.

Bejczy, A., Kim, W., and Venema, S., "The Phantom Robot: Predictive Displays for Teleoperation with Time Delay, " IEEE Robotics and Automation, 1990.

Boff, K., Kaufman, L., and Thomas, J., *Handbook of Perception and Human Performance*, John Wiley and Sons, New York, 1984

Brooks, T. L., "Telerobotic Response Requirements", Proc IEEE Int Conf on Systems, Man, and Cybernetics, Los Angeles, CA, November, pp. 113-120.

Brooks, Frederick P., Jr. Ming Ouh-Young, Michael Pique, John Hughes, Neela Srinivasan. 1988. "Using a Manipulator for Force Display in Molecular Docking." IEEE pub, CH2555-1. 1824-1829.

Brooks, Frederick P., Jr. Ming Ouh-Young, Batter, James J., Kilpatrick, Jerome P., 1990. "Project GROPE-Haptic Display for Scientific Visualization." Computer Graphics, Vol 24, 4,177-184

Brooks, V. B., *The Neural Basis of Motor Control.* Oxford University Press, New York, 1986.

Brooks, T., Ince, I. , and Lee, G., "Operator Vision Aids for Space Teleoperation Assembly and Servicing," Space Operations, Applications, and Research Symposium 1991.

Burdea, G., and Zhuang, J., "Dexterous Telerobotics with Force Feedback - an overview. Part I: Human Factors," Robotica vol 9, pp 171-178, 1991.

Burdea, G., and Zhuang, J., "Dexterous Telerobotics with Force Feedback - an overview. Part II: Control and implemenation," Robotica vol 9, pp 291-298, 1991.

Cannon D. J., "Point-and-Direct Telerobotics: Object Level Strategic Supervisory Control in Unstructured Interactive Human-Machine System Environments," Ph.D. Dissertation, Stanford University, June 1992.

Chin, K.P., *Stable teleoperation with Optimal Performance* Ph.D. dissertation, Dept. of Mech. Eng., M.I.T., Cambridge MA, 1991.

Colgate, J.E., Grafing, P.E., and Stanley, M.C. "Implementation of stiff virtual walls in force reflecting interfaces." Submitted to IEEE-VRAIS, Seattle WA, 1993.

Durlach, N.I., "Auditory Localization in teleoperator and virtual environment systems." Perception, 20, 543-554., 1991

Durlach, N.I., Rigopulos, A., Pang, X.D. "On the Externalization of Auditory Images", Presence, Volume 1, Number 2, Spring 1992.

Ellis, S.R., *Pictoral Communication in Virtual and Real Environments*. Taylor and Francis, New York, 1991.

Ericson, Mark A., McKinley, R.L., and Valencia, G. "Human Auditory Localization Performance in Azimuth.", Armstrong Laboratories Technical Report AL-TR-1991-0080. Jan 1991.

Fischer, R.D., Siva, K.V., "Specification and Design of Input Devices for Teleoperation", IEEE 1990.

Hartline, H.K., "Inhibition of activity of visual receptors by illuminating nearby retinal areas in the Limulus eye." Federation Proceedings, 8, 69., 1949

Helmholtz, H. C. F., "Physiological Optics", Vol II, trans. J. Southall. Rochester, N.Y.: Optical Society of America, 1924

Human-Machine Interfaces for Teleoperators and Virtual Environments, NASA Conference Publication 10071, pp. 108-113.

Grafing, P.E., *A Study of a Haptic Interface for a Virtual Environment*. M.S. thesis, Dept. of Mech. Eng., Northwestern University, Evanston IL, 1992.

Fasse, E.D., *On the Use and Representation of Sensory Information of the Arm by Robots and Humans*. Ph.D. dissertation, Dept. of Mech. Eng., M.I.T., Cambridge MA, 1992.

Ferrell, W. R. and Sheridan, T.B., "Supervisory control of remote manipulation." IEEE Spectrum 4, no. 10 October: 81-88.

Ferrel, W. R., "Remote Manipulation with Transmission Delay." M.I.T., NASA TN D-2665, 1965.

Ferrel, W.R., "Delayed Force Feedback." Human Factors, 8, 449-455. 1966.

Fischer, R.D., Siva, K.V., "Specification and Design of Input Devices for Teleoperation", IEEE 1990.

Fitts, P. M., "The Information Capacity of Human Motor Systems in Controlling the Amplitude of a Movement," Journal Experimental Psychology., vol 47, pp. 381-391, 1954.

Hannaford, B., "Stability and Performance Trade-Offs in Bilateral Teleoperation." IEEE Conf. on Robotics and Automation. 1989

Hashimoto, T. Sheridan, T.B., and Noyes, M.V. "Effects of Predictive Information in Teleoperation with Time Delay." Japan Journal of Ergonomics, 22,2,1986.

Held, R. and Durlach, N., "Telepresence, time delay and adaptation." In NASA Conference Publ. 10032. 1987, Also in Ellis, S.R., 1991.

Held, R., "Visual interfaces: the human perceiver." In Proceedings of Engineering Foundation Conference on Human-Machine Interfaces for Teleoperators and Virtual Environments (NASA SP), Santa Barbara, CA. Mar 4-9. 1990.

Hill, J.W., "Study of Modeling and Evaluationof Remote Manipulation Tasks with Force Feedback," Final Report, SRI Project 7696, JPL Contract 95-5170, 1979.

Jacobsen, S.C., Smith, F.M., Backman, D.K., and Iversen, E.K., "High Performance, High Dexterity, Force Reflective Teleoperator II," ANS Topical Meeting on Robotics and Remote Systems, New Mexico, 1991.

Jex, H.R., "Some criteria for teleoperators and virtual enviornments from experience with vehicle/operator simulation." *Human-Machine Interfaces for Teleoperators and Virtual Environments*, NASA Conference Publication 10071, pp. 42-47, 1991.

Johnsen. E.G. and Corliss W.R., "Teleoperators and Human Augmentation." NASA SP-5047. pp. 72. 1967.

Jones, L.A. and Hunter, I.W., "Human operator perception of mechanical variabels and their effects on tracking performance." In *Advances in Robotics*, ed. H. Kazerooni, Amer. Soc. Mech. Eng, New York, pp. 49-53, 1992.

Julian, R.G., Anderson, T.R., "Robotics Telepresence: Applications of Human Controlled Robots in Air Force Maintenance," Aerospace Simulation, ISBN 0-911801-28-6, 1988.

Kilpatrick, P.J., *The Use of a Kinethetic Supplement in an Interactive Graphics System.* Ph.D. dissertation, Dept. of Comp. Sci., University of North Carolina, Chapel Hill NC, 1976.

Kim, W., Tendick, F., and Stark, L., "Visual Enhancements in Pick and Place Tasks: Human Operators Controlling a Simulated Cylindrical Manipulator," IEEE Journal of Robotics and Automation, RA-3, 1987 (b).

Kotoku, T., "Predictive Display with Foprce Feedback and Its Application to Remote Manipulation System with Transmission Time Delay." Proc. IEEE International Conference on Intelligent Robotics and Systems, 1992.

Lettvin, J.Y., Maturana, H.R., McCulloch, W.S., & Pitts, W.H., "What the frog's eye tells the frog's brain." Proceedings of the Institute of Radio Engineers, 47, 1940-1951., 1959

Levine, M. and Shefer, J., *Fundementals of Sensation and Perception*, 2nd Edition, Pacific Grove CA: Brooks/Cole Publishing Co., 1981.

Loomis, J.M., "Distal Attribution and Presence.", Presence 1,1. 1993.

MB Associates., *Operations and Maintenance Manual. Manipulator Arms Systems*, April 1976.

Massimino, M. and Sheridan, T.B., "Variable force and visual feedback effects on teleoperator man-machine performance." In proceedings of NASA Conference on Space Telerobotics, Pasadena, CA., 1989.

McGovern, D. E., *Factors Effecting Control Allocation for Augmented Remote Manipulation.*, Doctoral Dissertation, Stanford University, 1975

Minsky, M., Ouh-Young, M., Steele, O., Brooks, F.P., Jr., and Behensky, M., "Feeling and seeing: issues in force display." *Computer Graphics, 24*, 235-243, 1990.

Norman, D.A., *The Design of Everyday Things.*, Doubleday Publishing, New York, 1988.

Noyes, M.V., and Sheridan, T.B., "A Novel Predictor for Telemanipulation Through a Time Delay." Proc. of 20th Annual Conference on Manual Control, NASA Ames Research Center, Moffet Field, CA, 1984.

Ouh-Young, M., Pique, M. Hughes, J., Srinivasan, N., and Brooks, F.P., Jr., "Using a manipulator for force display in molecular docking." *Proceedings, IEEE Int. Conf. on Robotics and Automation*, Philadelphia PA, pp. 1824-1829, 1988.

Ouh-Young, M. *Force Display in Molecular Docking*. Ph.D. dissertation, Dept. of Comp. Sci., University of North Carolina, Chapel Hill, NC, 1990.

Oxenberg, S., Landell, B. and Kan, E., "Geometric Database Maintenance Using CCTV Cameras and Overlay Graphics." SPIE, Vol. 1006, Space Station Automation IV, 1988.

Paul, R., Lindsay, T., and Sayers, C. "Time Delay Insensitive Teleoperation." Proc. IEEE International Conference on Intelligent Robotics and Systems, 1992.

Pepper, R.L. and Hightower, J.D., "Research Issues in Teleoperator Systems." 28th Annual Human Factors Society Meeting, San Antonio, Tx, 1984.

Pepper, R. L. and Kaomea, P. K., "Teleoperation: Telepresence and Performance Assessment. Annual International Ergonomics Society Meeting", Teleoperation and Control, July 1988.

Remis, S. J., and Repperger, D.W., "Quantifying Performance Degradation Due to the Human-Machine Interface of Telemanipulators." AAMRL-SR-90-510, December 1990.

Repperger, D.W., and Goodyear, C., "Active Controllers and Time Duration to Learn a Task." Conference on Maunual Control, 1985.

Repperger, D.W., Remis, S. J. , and Merrill, G., "Performance Measures of Teleoperation Using an Exoskeleton Device." Proc. 1990 IEEE Int. Conf. Robot. Automat., vol. 1, 552-557, 1990.

Rosenberg, L.B., *The use of Virtual Fixtures As Perceptual Overlays to Enhance Operator Performance in Remote Environments*. Technical Report AL-TR-1992-XXX, USAF Armstrong Laboratory, Wright-Patterson AFB OH, 1992. In publication.

Rosenberg, L.B., *The use of Virtual Fixtures to Enhance Performance in Teleoperation with Time Delay*. Technical Report AL-TR-1993-XXX, USAF Armstrong Laboratory, Wright-Patterson AFB OH, 1992. In publication.

Rosenberg, L.B., *Perceptual Design of a Virtual Rigid Surface Contact*. Technical Report AL-TR-1993-XXX, USAF Armstrong Laboratory, Wright-Patterson AFB OH, 1993. In publication.

Rosenberg, L.B., "The Use of Virtual Fixtures to Enhance Operator Performance in Remote Environments" SPIE Telemanipulator Technology, Boston, MA 1993.

Rosenberg, L.B., "The Use of Virtual Fixtures to Enhance Teleoperator Performance in Remote Environments with Time Delay." Proceedings, ASME Winter Annual Meeting on Haptic Interfaces for Virtual Environments and Teleoperator Systems, New Orleans, Louisiana 1993

Rosenberg, L.B., "Virtual Fixtures: Perceptual Tools for Telerobotic Manipulation." Proceedings IEEE Virtual Reality Annual International Symposium, 1993

Rosenberg, L.B., Adelstein, B.D., "The Perceptual Decomposition of Virtual Haptic Surfaces." Proceedings of IEEE Symposium in Research Frontiers in Virtual Reality. 1993.

Russell, R. and D'Arcy A., "Video-Based Satellite Attitude Determination.' SPIE, Vol. 729, Space Station Automation II, 1986.

Sheridan, T., 'Human Supervisiory Control of Robot Systems." IEEE Robotics and Automation, Vol 2, 1986.

Sheridan, T., "Telerobotics", Automatica, 25(4), 1989.

Sheridan, T. B. and Ferrell, W.R., *Man-Machine Systems*. The MIT Press, 1981.

Sheridan, T.B., *Telerobotics, Automation, and Human Supervisory Control.*, MIT Press, Cambridge Mass, 1992. pp 208-211.

Sheridan, Thomas B., "Defining our Terms.", Presence 1,2. 272-274. 1993

Slater, M. and M. Usoh., " Presence in Immersive Virtual Environments", Proceedings VRAIS 93, . Seattle Washington, September 1993.

Spain, Edward, H., *Peg-In-Hole Taskboard Documentation*. Final Report for Task FQ7624-88-00024, March 1989.

Stark, L., 'Telerobotics: Display, Control, and Communications Problems.' IEEE Robotics and Automation 1989.

Tan, H.Z., Pang, X.D., and Durlach, N.I., "Manual resolution of length, force, and compliance." In *Advances in Robotics*, ed. H. Kazerooni, Amer. Soc. Mech. Eng. New York, pp. 13-18, 1992.

Thompson, D. A., "The Development of a Six Degree of Freedom Robot Evaluation Test. In Proceedings of 13th Annual Conference on Manual Control." MIT, Cambridge MA, 1977

Wallach, H., "The role of head movements and vestibular and visual cues in sound localization." Journal of Experimental Psychology, 27, 339-368. 1940.

White, B. W., "Perceptual findings with a vision substitution system." IEEE Transactions of Man-Machine Systems, MMS-11, 54-59. 1970.

White, B.W., Saunders, F. A., Scadden, L., Bach-y-Rita, P., & Collins, C.C., "Seeing with the skin.", Perception & Psychophysics, 7, 23-27. 1970.

Winey, C.M., III., *Computer Simulated Visual and tactile Feedback as an Aid to Manipulator and Vehicle Control*. S.M. thesis, Dept. of Mech. Eng., M.I.T., Cambridge, MA, 1981.

Young, P.T., "The role of head movements in auditory localization." Journal of Experimental Psychology, 14, 95-124., 1931.

The Virtual Fixtures project was funded by the Air Force Office of Scientific Research (AFOSR) and is therefore entered into the public domain. The research was conducted at Stanford University, Air Force Research Laboratory (Armstrong Lab), and NASA Ames Research Center. The objective of the Virtual Fixtures platform was to research whether it was possible to incorporate interactive virtual objects (called "virtual fixtures") into a user's perception of the real world. This concept later became commonly referred to as augmented reality (AR) or mixed reality (MR). Also see USAF Technical Reports AL/CF-TR-94-0089 and 94-0139.

"Virtual fixtures": Perceptual overlays enhance operator performance in telepresence tasks

Rosenberg, Louis Barry, Ph.D.

Stanford University, 1994

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106